

**Probabilistic Horn abduction  
and  
Bayesian networks**

by  
David Poole

Technical Report 92-2  
January 1992

---

Department of Computer Science  
University of British Columbia  
Rm 333 - 6356 Agricultural Road  
Vancouver, B.C.  
CANADA V6T 1Z2



# Probabilistic Horn abduction and Bayesian networks

David Poole

Department of Computer Science,  
University of British Columbia,  
Vancouver, B.C., Canada V6T 1Z2

poole@cs.ubc.ca

telephone: (604) 822 6254

fax: (604) 822 5485

January 13, 1992

## Abstract

This paper presents a simple framework for Horn-clause abduction, with probabilities associated with hypotheses. The framework incorporates some assumptions about the rule base and some independence assumptions amongst hypotheses. It is shown how any probabilistic knowledge representable in a discrete Bayesian belief network can be represented in this framework. The main contribution is in finding a relationship between logical and probabilistic notions of evidential reasoning. This provides a useful representation language in its own right, providing a compromise between heuristic and epistemic adequacy. It also shows how Bayesian networks can be extended beyond a propositional language, and shows a relationship between probability and argument based systems.



## 1 Introduction

Probabilistic Horn Abduction [43, 42] is a framework for logic-based abduction that incorporates probabilities with assumptions. This is being used as a framework for diagnosis [43] that incorporates both pure Prolog and Bayesian Networks as special cases [42]. This paper expands on [43, 42] and develops the formal underpinnings of probabilistic Horn abduction, shows the strong relationships to other formalisms and argues that it is a good representation language in its own right. It can be motivated in a number of different ways:

Determining what is in a system from observations (diagnosis and recognition) is an important part of AI. There have been many logic-based proposals as to what is a *diagnosis* [16, 51, 13, 40, 12]. One problem with all of these proposals is that for any diagnostic problem of a reasonable size there are far too many logical possibilities to handle. For example, when considering fault models [14, 40], there is almost always an exponential number of logical possibilities (e.g., each component could be in one of its normal states or in the unknown state). For practical problems, we find that many of the logically possible diagnoses are so unlikely that it is not worth considering them. There is a problem, however, in removing the unlikely possibilities a priori (those with a low prior probability): it may happen that the unlikely occurrence is the actual truth in the world. Analysis of the combinatorial explosions would however tend to suggest that we need to take into account probabilities of the diagnoses [13, 35, 31], and not even generate the unlikely diagnoses (i.e., those with a low posterior probability).

In a different strand of research, Bayesian networks [34], have proven to be a good representation for the sort of probabilistic independence found in many domains. While the independence of Bayesian networks has been expressed in logic (e.g., [3]), there has not been a mapping between logical specifications of knowledge and Bayes net representations, where the logic is not at the meta-level to the probabilistic knowledge. This paper describes what could be termed as *the* logic of discrete Bayesian Networks, where we use the logic to express the object level knowledge of the Bayesian network. In the uncertainty community there has also been a need to extend Bayesian networks to beyond a propositional language [5, 6, 20]. Probabilistic Horn abduction is naturally non-propositional, and provides a natural extension of Bayesian networks to a non-propositional language. The work presented in this paper should be contrasted with other attempts to combine logic and

probability in very powerful languages (e.g., [3]). We are trying to find the simplest language that is useful for our purposes, rather than combine many different features onto one framework. Our goal in this research is to find the simplest logic that is useful for our purposes.

The representation proposed in this paper is interesting in its own right as a compromise between epistemic and heuristic adequacy [30]. It extends pure Prolog in a simple way to include probabilities. While all of the hypotheses are independent, by inventing new hypotheses, we can represent any probabilistic dependency. This simplicity allows us to experiment with a minimalist representation and only extend it when we need to. It is interesting to see how far we can go with a very simple representation language, only adding to it when it fails to do what we want to do.

### 1.1 A Motivating Example

Before we present the language and the assumptions behind the representation, we first give an example to show what sorts of things we can represent.

The example is based on the three cascaded inverters of [14]. Figure 1 shows the connections of the inverters.




Figure 1: Three cascaded inverters.

The language is an extension of pure Prolog (without negation as failure). We can write Prolog-like definite clauses to represent the general knowledge of the domain, and the specific instances known about the particular configuration:

$$\begin{aligned} \text{val}(\text{out}(G), \text{on}, T) &\leftarrow \text{ok}(G) \wedge \text{val}(\text{in}(G), \text{off}, T). \\ \text{val}(\text{out}(G), \text{off}, T) &\leftarrow \text{ok}(G) \wedge \text{val}(\text{in}(G), \text{on}, T). \end{aligned}$$

$val(out(G), V, T) \leftarrow shorted(G) \wedge val(in(G), V, T).$   
 $val(out(G), off, T) \leftarrow blown(G).$   
 $val(in(G), V, T) \leftarrow conn(G1, G) \wedge val(out(G1), V, T).$   
 $conn(i1, i2).$   
 $conn(i2, i3).$

Here  $val(P, V, T)$  means that port  $P$  has value  $V$  at time  $T$ ;  $in(G)$  is the input port of gate  $G$  and  $out(G)$  is the output port of gate  $G$ .  $ok(G)$  means  $G$  is working properly;  $shorted(G)$  means that  $G$  is shorted and acts as a wire; and  $blown(G)$  means that  $G$  always outputs the value  $off$ .

The language also has a "disjoint declaration" that defines a set of disjoint and covering hypotheses that have probabilities associated with them:

$disjoint([ok(G) : 0.95, shorted(G) : 0.03, blown(G) : 0.02]).$   
 $disjoint([val(in(i1), on, T) : 0.5, val(in(i1), off, T) : 0.5]).$

The first gives the prior probabilities of the states of gates. The second gives the prior probabilities of the inputs to the first gate. In our language, we have stated that the gates break independently, and that the values of the input to gate  $i1$  is independent of the states of the system.

The sorts of things that we can ask, and for which we have given enough information to compute include:

- What is the probability that gate  $i2$  is ok given that the input to  $i1$  is  $off$  and the output of  $i3$  is  $off$  at time  $t_1$ ?

$$P(ok(i2) | val(in(i1), off, t_1) \wedge val(out(i3), off, t_1))$$

(The answer is 0.76).

- If the input of  $i1$  were  $on$  what is the probability that the output of  $i3$  will be  $off$ ?

$$P(val(out(i3), off, t_1) | val(in(i1), on, t_1))$$

(The answer is 0.899).

- What was the probability that the input to  $i1$  was  $on$  at time  $t_2$  given that the output to  $i2$  was  $off$  at time  $t_2$  and given that the output of  $i3$  was  $off$  and the input of  $i1$  was  $off$  at time  $t_1$ ?

$$P(val(in(i1), on, t_2) | val(out(i2), off, t_2) \wedge val(out(i3), off, t_1) \wedge val(in(i1), off, t_1))$$

(the answer is 0.55).

Each of these answers is computed in terms of explanations, namely arguments with premises from the possible hypotheses. We assume independence amongst the hypotheses so that the prior probability of an explanation is obtained by multiplying the probabilities of the hypotheses in the explanation.

For example, the explanations of the observation  $val(in(i1), off, t_1) \wedge val(out(i3), off, t_1)$ , together with their prior probability are:

Explanation:  $[val(in(i1), off, t_1), ok(i3), ok(i2), shorted(i1))]$   
Prior = 0.01354

Explanation:  $[val(in(i1), off, t_1), ok(i3), shorted(i2), ok(i1))]$   
Prior = 0.01354

Explanation:  $[val(in(i1), off, t_1), shorted(i3), ok(i2), ok(i1))]$   
Prior = 0.01358

Explanation:  $[val(in(i1), off, t_1), blown(i3)]$   
Prior = 0.01

Explanation:  $[val(in(i1), off, t_1), ok(i3), ok(i2), blown(i1))]$   
Prior = 0.009025

Explanation:  $[val(in(i1), off, t_1), shorted(i3), blown(i2))]$   
Prior = 0.0003

Explanation:  $[val(in(i1), off, t_1), shorted(i3), shorted(i2), shorted(i1))]$   
Prior = 1.35e-5

Explanation:  $[val(in(i1), off, t_1), shorted(i3), shorted(i2), blown(i1))]$   
Prior = 9.0e-6

By the way the knowledge base was constructed, these explanations are disjoint and covering, and so we can compute the prior probability of

$$val(in(i1), off, t_1) \wedge val(out(i3), off, t_1))$$

by summing the probabilities of these explanations, which here is 0.05996.



## 2 Probabilistic Horn Abduction

In this section we develop the language, the assumptions behind and a “semantics” of the language in terms of abduction. The language is designed to be usable and does not allow us to state what cannot be computed in a straight forward manner.

The initial language is translated into an abductive framework with a number of assumptions about the knowledge base. The appendix gives a more formal model-theoretic semantics and demonstrates the equivalence between the two.

### 2.1 The Probabilistic Horn abduction language

Our language uses the Prolog conventions:

**Definition 2.1** A term is either a variable (starting with an upper case letter), a constant (starting with a lower case letter) or is of the form  $f(t_1, \dots, t_n)$  where  $f$  is a function symbol (starting with a lower case letter) and each  $t_i$  is a term. An atomic symbol (atom) is of the form  $p$  or  $p(t_1, \dots, t_n)$  where  $p$  is a predicate symbol (starting with a lower case letter) and each  $t_i$  is a term.

**Definition 2.2** A definite clause is of the form:  $a$ . or  $a \leftarrow a_1 \wedge \dots \wedge a_n$ . where  $a$  and each  $a_i$  are atomic symbols.

**Definition 2.3** A disjoint declaration is of the form

$$\text{disjoint}([h_1 : p_1, \dots, h_n : p_n]).$$

where the  $h_i$  are atoms, and the  $p_i$  are real numbers  $0 \leq p_i \leq 1$  such that  $p_1 + \dots + p_n = 1$ . Any variable appearing in one  $h_i$  must appear in all of the  $h_j$  (i.e., the  $h_i$  share the same variables). The  $h_i$  will be referred to as hypotheses or assumables.

**Definition 2.4** A probabilistic Horn abduction theory (which will be referred to as a “theory”) is a collection of definite clauses and disjoint declarations such that if a ground atom  $h$  is an instance of a hypothesis in one disjoint declaration, then it is not an instance of another hypothesis in any of the disjoint declarations.

Given theory  $T$ , we define

$F_T$  the **facts**, is the set of definite clauses in  $T$  together with the clauses of the form

$$\text{false} \leftarrow h_i \wedge h_j$$

where  $h_i$  and  $h_j$  both appear in the same disjoint declaration in  $T$ , and  $i \neq j$ . Let  $F'_T$  be the set of ground instances of elements of  $F_T$ .

$H_T$  the **hypotheses**, the set of  $h_i$  that appears in some disjoint declaration in  $T$ . Let  $H'_T$  be the set of ground instances of elements of  $H_T$ .

$P_T$  is a function  $H'_T \mapsto [0, 1]$ .  $P(h'_i) = p_i$  where  $h'_i$  is a ground instance of hypothesis  $h_i$ , and  $h_i : p_i$  is in a disjoint declaration in  $T$ .  $P(h'_i)$  will be the prior probability of  $h'_i$ .

Where  $T$  is understood from context, we omit the subscript.

The disjoint declarations allow a very restricted form of integrity constraints [23]. It allow binary integrity constraints (the conjunction of two hypotheses is false) such that the ground instances of hypotheses form mutually exclusive and covering groupings that correspond to random variables.

A theory will define a set of **represented atoms** that are a subset of the atoms of  $T$ . The represented atoms will often be not listed explicitly, but will be left implicit (they will typically be instances of hypotheses and heads of clauses).

The represented atoms are those about which the theory can answer questions. Questions about atoms not in the represented atoms will be beyond the scope of the theory. The theory is not expected to be able to answer queries outside of its scope.

## 2.2 Abduction

We first give the language an abductive characterisation, using the normal definition of the definite clauses. This is used to make explicit our assumptions and to build the theory in a natural manner. In Appendix A, we give a model theoretic characterisation that incorporates our assumptions, and show the equivalence of the formulations.

The formulation of abduction used is a simplified form [17] of Theorist [45, 38]. It is simplified in being restricted to Horn clauses. This can also

be seen as a generalisation of an ATMS (with predefined nogoods) [52] to be non-propositional<sup>1</sup>.

An abductive scheme is a pair  $\langle F, H \rangle$  where

$F$  is a set of Horn clauses. Variables in  $F$  are implicitly universally quantified. Let  $F'$  be the set of ground instances of elements of  $F$ .

$H$  is a set of (possibly open) atoms, called the “assumables” or the “possible hypotheses”. Let  $H'$  be the set of ground instances of elements of  $H$ .

**Definition 2.5** [45, 37] If  $g$  is a closed formula, an **explanation** of  $g$  from  $\langle F, H \rangle$  is a set  $D$  of elements of  $H'$  such that

- $F \cup D \models g$  and
- $F \cup D \not\models \text{false}$ .

The first condition says that  $D$  is sufficient to imply  $g$ , and the second says that  $D$  is possible.

**Definition 2.6** A **minimal explanation** of  $g$  is an explanation of  $g$  such that no strict subset is an explanation of  $g$ .

### 2.3 Assumptions about the rule base

In order to be able to simply interpret our rules probabilistically we make some assumptions about the rules and some probabilistic independence assumptions about the hypotheses.

The first assumption is syntactic, about the relationship between hypotheses and rules:

**Assumption 2.7** There are no rules in  $F$  whose head unifies with a member of  $H$ .

---

<sup>1</sup>A main difference is in the philosophy of use. We assume that the Horn clauses are representing the object level knowledge, rather than, as in an ATMS, acting as a back end of a problem solver [11].

This does not seem to be a very severe restriction, in practice. It says that we do not want rules to imply a hypothesis. Presumably, if we had rules for possible hypothesis  $h$ , the only reason that we would want to make  $h$  a hypothesis is if it was possible that  $h$  just happens to be true (without any other “cause”). In this case we can replace  $h$  in  $H$  by the hypothesis  $h\_happens\_to\_be\_true$  and add the rule

$$h \leftarrow h\_happens\_to\_be\_true$$

**Assumption 2.8** (acyclicity) If  $F'$  is the set of ground instances of elements of  $F$ , then it is possible to assign a natural number to every ground atom such that for every rule in  $F'$  the atoms in the body of the rule are strictly less than the atom in the head.

This assumption is described as natural by Apt and Bezem [1] (they, however, also include negation as failure, which we do not (but see section 2.5.3)). It is a generalisation of the hierarchical constraint of Clark [8]. It implies that there are no infinite chains when backchaining from any ground goal. This does not restrict recursion, but does mean that all recursion must be well founded.

These assumptions are made implicitly in [38], are explicit in [9], but are relaxed in [22].

When using abduction we often assume that the explanations are covering. This can be a valid assumption if we have anticipated all eventualities, and the observations are within the domain of the expected observations (usually if this assumption is violated there are no explanations). This is also supported by recent attempts at a completion semantics for abduction [38, 9, 22]. The results show how abduction can be considered as deduction on the “closure” of the knowledge base that includes statements that the given causes are the only causes. We make this assumption explicit here:

**Assumption 2.9** The rules in  $F'$  for every ground non-hypothesis represented atom are covering.

That is, if the rules for  $a$  in  $F'$  are

$$\begin{aligned} a &\leftarrow B_1 \\ a &\leftarrow B_2 \\ &\vdots \\ a &\leftarrow B_m \end{aligned}$$

if  $a$  is true, one of the  $B_i$  is true. The *completion* of  $a$  is

$$a \equiv B_1 \vee \dots \vee B_n$$

Thus assumption 2.9 says that Clark's completion [8] is valid for every non-assumable.

If the rules for  $a$  are not covering, then we create a new hypothesis *a\_is\_true\_for\_some\_other\_reason* and add the rule

$$a \leftarrow a\_is\_true\_for\_some\_other\_reason.$$

**Lemma 2.10** [38, 9] *Under assumptions 2.7, 2.8 and 2.9, if  $\text{expl}(g, T)$  is the set of all minimal explanations of  $g$  from  $\langle F_T, H_T \rangle$ , and  $\text{comp}(T)$  is  $F_T$  augmented with the completion of every ground instance of every non-assumable, then*

$$\text{comp}(T) \models g \equiv \bigvee_{e_i \in \text{expl}(g, T)} e_i$$

The next assumption has to do with the status of explanations

**Assumption 2.11** The bodies of the rules in  $F'$  for an atom are mutually exclusive.

Given the above rules for  $a$  this means that  $B_i \wedge B_j$  is always false for each  $i \neq j$ . To ensure that this assumption holds we can add extra conditions to the rules. See section 5.

Note that, whereas assumptions 2.7 and 2.8 are syntactic assumptions about the theory that can be automatically checked, assumptions 2.9 and 2.11 are statements about the world, and not about the knowledge base<sup>2</sup>. We do not require  $F_T \models \neg(B_i \wedge B_j)$ . The language is not powerful enough to state such constraints. For example, it may be the case that the value on some wire is functional (the value cannot be both on and off at the same time). We cannot state this in our language. This is a deliberate design decision to make the language as simple as possible. See section 2.5 for the rationale behind the design decisions.

<sup>2</sup>There is, however, a syntactic condition that can be used to check whether assumption 2.11 has been violated. This is when we can derive the bodies of two rules for an atom from a set of assumptions that are not inconsistent.

**Lemma 2.12** *Under assumptions 2.7 and 2.11, minimal explanations of atoms or of conjunctions of atoms are mutually exclusive (no two explanations can both be true).*

Lemma 2.12 does not hold for arbitrary formulae. In particular, the minimal explanations of a disjunction are not necessarily disjoint.

## 2.4 Probabilities

Associated with each possible hypothesis is a prior probability. We use this prior probability to compute arbitrary probabilities.

The following is a corollary of lemmata 2.10 and 2.12

**Lemma 2.13** *Under assumptions 2.7, 2.8, 2.9 and 2.11, if  $\text{expl}(g, T)$  is the set of minimal explanations of conjunction of atoms  $g$  from probabilistic Horn abduction theory  $T$ :*

$$\begin{aligned} P(g) &= P\left(\bigvee_{e_i \in \text{expl}(g, T)} e_i\right) \\ &= \sum_{e_i \in \text{expl}(g, T)} P(e_i) \end{aligned}$$

Thus to compute the prior probability of any  $g$  we sum the probabilities of the explanations of  $g$ .

To compute arbitrary conditional probabilities, we use the definition of conditional probability:

$$P(\alpha|\beta) = \frac{P(\alpha \wedge \beta)}{P(\beta)}$$

To find arbitrary conditional probabilities  $P(\alpha|\beta)$ , we find  $P(\beta)$ , which is the sum of the explanations of  $\beta$ . To compute the probability  $P(\alpha \wedge \beta)$ , we sum over the explanations of  $\alpha \wedge \beta$ . Note that the explanations of  $\alpha \wedge \beta$  are also explanations of  $\beta$ . We can find the explanations of  $\alpha \wedge \beta$  by explaining  $\alpha$  from the explanations of  $\beta$ . Thus arbitrary conditional probabilities can be computed from summing the prior probabilities of explanations.

It remains only to compute the prior probability of an explanation  $D$  of  $g$ . We assume that logical dependencies impose the only statistical dependencies on the hypotheses. In particular we assume:

**Assumption 2.14** Ground instances of hypotheses that are consistent (with  $F_T$ ) are probabilistically independent.

**Example 2.15** If we have the disjoint declarations

$$\begin{aligned} & disjoint([p(X) : 0.4, q(X) : 0.6]). \\ & disjoint([r(Y) : 0.01, s(Y) : 0.99]). \end{aligned}$$

then  $p(t)$  is independent of  $r(u)$  for all ground terms  $t$  and  $u$ .  $p(t)$  is independent of  $q(u)$  for all different ground terms  $u$  and  $t$ .  $p(t)$  is independent of  $p(u)$  for all different ground terms  $u$  and  $t$ .

Thus  $p(a)$  is dependent on  $q(a)$  (they are exclusive), but  $p(a)$  is independent of  $q(b)$ .  $p(a)$  is also independent of  $p(b)$ .

Under assumption 2.14, if  $D = \{h_1, \dots, h_n\}$  is a minimal explanation, then

$$P(h_1 \wedge \dots \wedge h_n) = \prod_{i=1}^n P(h_i)$$

To compute the prior of the minimal explanation we multiply the priors of the hypotheses.

Assumption 2.14 implies the unique names assumption if there are parametrized hypotheses. If  $h(X)$  is a parametrized hypothesis with each instance having probability  $p$ , and  $t_1$  and  $t_2$  are different ground terms, assumption 2.14 implies  $P(h(t_1) \wedge h(t_2)) = p^2$ . If  $t_1 = t_2$ , then  $h(t_1) \wedge h(t_2) \equiv h(t_1)$ , but  $P(h(t_1)) = p$ , a contradiction to the fact that probabilities are a measure over propositions, and that logically equivalent terms should have the same probability [34]. Thus, we are assuming that  $t_1 \neq t_2$  for different terms  $t_1$  and  $t_2$ . This assumption is the unique names assumption [50]. Note that it is for the probabilistic calculation that we are making the unique names assumption.

Appendix A gives the formal semantics for probabilistic Horn abduction, and justifies, in another way, the results of this section.

## 2.5 Rationale for the design choices

The language presented so far is quite weak in some respects. In this section we discuss why the language is as it is. The general theme is that the language

is a simple extension to pure Prolog that lets us consistently interpret the numbers on hypotheses as probabilities. We have disallowed anything that will make this interpretation difficult. For example, we have not allowed the logic to be expressive enough to be able to prove that there is a dependency amongst the hypotheses beyond the disjointness of our random variables.

This simplicity makes the language semantically transparent, and allows for simple implementations. It is still powerful enough to express many of the causal and probabilistic interactions that we want to express. This work should be seen as an exercise in minimalist representations — we try to understand the limitations of very restricted languages and only add extra power if we can show we cannot do what we want with the tools available to us.

### 2.5.1 Language for specifying random variables

The first thing to notice is that we only allow a very restricted and stylised form of integrity constraints to be specified by the use of the disjoint declaration. This is in contrast to earlier versions [43, 42] where we allowed arbitrary integrity constraints. While this more expressive language allows us to represent what the current version allows, it also lets us represent what we cannot interpret (both in terms of interpreting the numbers as probabilities and in terms of making the proof procedures more complicated without providing visible advantage).

For example, if  $\{a, na\}$  and  $\{b, nb\}$  each form a disjoint and covering sets of propositions (random variables), then we cannot treat these as independent if we can state  $false \leftarrow a \wedge b$ . All we know here is that the variables are not independent — there are only ad hoc methods to allow us to provide the joint distribution. The current formulation does not give the power to state such constraints. The logical formulae provide no constraints on the hypotheses beyond the disjointedness of the values in the disjoint declaration.

A second, but related, problem with general integrity constraints<sup>3</sup> has to do with making implicit assumptions by the use of integrity constraints. For example, if we have  $false \leftarrow a \wedge b$ , when we are using  $a$ , we are implicitly assuming the negation of  $b$  and should pay the cost (in terms of making any explanation containing  $a$  less likely). This occurs at an extreme level when

<sup>3</sup>This was pointed out to me by Mark Wallace of ECRC.



we have  $\{b, nb\}$  disjoint hypotheses whose probability sums to 1, and have  $false \leftarrow a \wedge b$  and  $false \leftarrow a \wedge nb$ . Here it should follow that  $a$  cannot occur (or at least with probability zero) and should be pruned from other explanations.

We also made sure that all atoms in disjoint declarations share the same variables. To see the problem with not requiring this, consider the (illegal) declaration

$$disjoint([p : 0.7, q(X) : 0.3]).$$

Given this declaration,  $p$  would be disjoint with  $q(a)$  and  $p$  disjoint with  $q(b)$ . This would then place a dependence between  $q(a)$  and  $q(b)$ . Given that the hypotheses are covering, if  $q(a)$  is true, then  $p$  is false, and so  $q(b)$  is true. Similarly if  $q(a)$  is false,  $p$  is true and  $q(b)$  is false. Thus all instances of  $q(X)$  would always have the same truth value. Then we may as well remove the variable (as the truth does not depend on the value of the variable). We have restricted the variables in disjoint assertions in order to avoid such tricky interactions that the user of the system may not be aware of, and so that an implementation does not need to look for them.

The language is also not powerful enough to state the constraints on the legal input. When we wanted the bodies of the rules to be disjoint, we did not require that we could prove the disjointedness of the bodies. We only required that they be disjoint in the domain under consideration. Similarly, we require the rules to be covering, but cannot check this. This should not be seen as a defect in this language — there will always be true things that the language is too weak to state (e.g., the finiteness of integers cannot be stated in the first order predicate calculus). This restriction was a conscious decision to allow us to build efficient implementations, and to avoid difficult to interpret statements as described above.

### 2.5.2 Assumptions concerning the knowledge base

Assumption 2.7 is important to ensure that we can treat the hypotheses as independent. If we could prove some hypothesis (based on other assumptions) it would not be consistent that the hypotheses are independent. It is also important to ensure minimal explanations are disjoint. If we have  $a$  and  $b$  as hypotheses, and have  $b \leftarrow a$  as well as  $c \leftarrow b$ . There are two minimal explanations of  $c$ , namely  $\{a\}$  and  $\{b\}$ . These cannot be disjoint as one implies the other (and are equivalent under the Covering Assumption 2.9).

To see the importance of the acyclicity assumption 2.8 consider the facts  $F = \{a \leftarrow a \wedge b, a \leftarrow c\}$ , and possible hypotheses  $H = \{b, c\}$ . Assumption 2.8 is violated in this example. Here there is one explanation for  $a$ , namely  $c$ , but under assumption 2.9, we can only prove  $a \Rightarrow b \vee c$ , and not  $a \Rightarrow c$ . This violates the conclusion of Lemma 2.10.

Assumptions 2.9 and 2.11 are needed so that we can have disjoint and covering hypotheses. This means that we just sum the probabilities of the hypotheses. This seems like a much more natural assumption, and is much easier to enforce than, for example, the assumption that the bodies of the rules are independent given the head [24, 48].

### 2.5.3 Negation

The language provided does not have explicit negation. It does, however, give us an implicit negation.

For example, suppose we have the theory

$$\begin{aligned} a &\leftarrow h_1. \\ b &\leftarrow h_2. \\ \text{disjoint}([h_1 : p_1, h_2 : 1 - p_1]). \end{aligned}$$

Under assumption 2.9, we have  $a \equiv h_1$  and  $b \equiv h_2$ . The disjoint declaration essentially tells us that  $h_1 \equiv \neg h_2$ . Thus we have  $a \equiv \neg b$ .

Thus, although we cannot state negation, we can interpret one atom as being the negation of another. The reason that we don't want to have explicit negation is that this would allow the logic to imply a dependence amongst variables that violates the independence assumption 2.14. Once this occurs, it is more difficult to interpret the probabilities. For example,  $a \leftarrow h_1$  and  $\neg a \leftarrow h_3$  places a dependence on  $h_1$  and  $h_3$ .

For each atom we can create its negation. For each atom  $a$ ,  $\bar{a}$  is another atom which we interpret as the negation of  $a$ . Syntactically,  $\bar{a}$  will be just another atom.

If we have a disjoint declaration

$$\text{disjoint}([h_1 : p_1, h_2 : p_2, \dots, h_n : p_n]).$$

we can create the negation of any hypothesis, say  $h_1$ , by using

$$\bar{h}_1 \leftarrow h_2$$

$$\begin{array}{c} \vdots \\ \bar{h}_1 \leftarrow h_n \end{array}$$

The other hypotheses can be negated analogously.

If we have rules for  $a$

$$\begin{array}{c} a \leftarrow b_{11} \wedge \cdots \wedge b_{1n_1} \\ \vdots \\ a \leftarrow b_{k1} \wedge \cdots \wedge b_{kn_k} \end{array}$$

we can define atom  $r_i$  to correspond to the  $i$ -th body, and have rules  $a \leftarrow r_1$  through  $a \leftarrow r_k$ . There are also  $k$  rules of the form  $r_i \leftarrow b_{i1} \wedge \cdots \wedge b_{in_i}$ .

We can define the atom that is the negative of  $a$ , written here as  $\bar{a}$  as

$$\bar{a} \leftarrow \bar{r}_1 \wedge \cdots \wedge \bar{r}_k$$

$\bar{r}_i$  is defined as

$$\begin{array}{c} \bar{r}_i \leftarrow \bar{b}_{i1} \\ \vdots \\ \bar{r}_i \leftarrow \bar{b}_{in_i} \end{array}$$

These definitions will be well grounded by the acyclicity of the rule base.

Note that just because some atom  $a$  is represented, it does not mean that the negation of  $a$  need be represented. There are however, some cases for which it is useful to create the negation of atoms (see section 5.1).

This negation is closely related to that of Barbuti et. al. [4], which is used for negation as failure. The combination of the two ideas is, however, beyond the scope of this paper.

### 3 Representing Bayesian networks

In this section we give the relationship between Bayesian networks and our probabilistic Horn abduction. We show how any probabilistic knowledge that can be represented in a discrete Bayesian network, can be represented in our

formalism. We also demonstrate the alternate, namely that any propositional probabilistic Horn abduction theory is equivalent to a Bayes net.

A *Bayesian network* [34] is a directed acyclic network where the nodes represent random variables, and the arcs represent a directly influencing relation. We will use the term "RV" to mean random variable so as to avoid confusion with the Prolog-style variable. An arc from RV  $b$  to RV  $a$  represents the fact that RV  $b$  *directly influences* RV  $a$ ;  $b$  is said to be a parent of  $a$ . The relation *influences* is the transitive closure of the directly influences relation (corresponding to the ancestor relation).

Suppose we have a Bayesian network with random variables  $a_1, \dots, a_n$ , such that random variable  $a_i$  can have values  $v_{i,1}, \dots, v_{i,r_i}$ . We represent random variable  $a_i$  having value  $v_{i,j}$  as the proposition  $a_i(v_{i,j})$ .

Suppose RV  $a_i$  is directly influenced by RVs  $\Pi_{a_i} = \{a_{i_1}, \dots, a_{i_{n_i}}\}$  (the "parents" of  $a_i$ ) in a Bayesian network. The independence assumption embedded in a Bayesian Network [34] is given by

$$P(a_i | \Pi_{a_i} \wedge v) = P(a_i | \Pi_{a_i})$$

where  $v$  is a RV (or conjunction of RVs) such that  $a_i$  does not influence  $v$  (or any conjunct in  $v$ ).

The formal definition of a Bayes net is often given in terms of the joint distribution of all of the RVs:

$$P(a_1, \dots, a_n) = \prod_{i=1}^n P(a_i | \Pi_{a_i})$$

A Bayesian network is represented by Probabilistic Horn abduction rules that relates a RV with its parents:

$$a_i(V) \leftarrow a_{i_1}(V_1) \wedge \dots \wedge a_{i_{n_i}}(V_{n_i}) \wedge c_{\cdot}a_i(V, V_1, \dots, V_{n_i})$$

The intended interpretation of  $c_{\cdot}a_i(V, V_1, \dots, V_{n_i})$  is that  $a_i$  has value  $V$  because  $a_{i_1}$  has value  $V_1, \dots$ , and  $a_{i_{n_i}}$  has value  $V_{n_i}$ .

Associated with the Bayesian network is a contingency table which gives the marginal probabilities of the values of  $a$  depending on the values of  $\Pi_{a_i} = \{a_{i_1}, \dots, a_{i_{n_i}}\}$ . This will consist of probabilities of the form

$$P(a_i = v_{i,j} | a_{i_1} = v_1, \dots, a_{i_{n_i}} = v_{n_i}) = p_j$$

such that

$$\forall v_1, \dots, v_{n_i} \left( \sum_{j=1}^{r_i} P(a_i = v_{i,j} | a_{i_1} = v_1, \dots, a_{i_{n_i}} = v_{n_i}) \right) = 1$$

where  $v_{i,1}, \dots, v_{i,r_i}$  are the possible values for RV  $a_i$ . This is translated into assertions

$$\text{disjoint}([c\_a(v_{i,1}, v_1, v_2, \dots, v_{n_i}) : p_1, \dots, c\_a(v_{i,r_i}, v_1, v_2, \dots, v_{n_i}) : p_{r_i}])$$

**Example 3.1** Consider a representation of the Bayesian network of Figure 3.1, with the following conditional probability distributions:

$$\begin{aligned} P(\text{fire}) &= 0.01 \\ P(\text{smoke}|\text{fire}) &= 0.9 \\ P(\text{smoke}|\neg\text{fire}) &= 0.01 \\ P(\text{tampering}) &= 0.02 \\ P(\text{alarm}|\text{fire} \wedge \text{tampering}) &= 0.5 \\ P(\text{alarm}|\text{fire} \wedge \neg\text{tampering}) &= 0.99 \\ P(\text{alarm}|\neg\text{fire} \wedge \text{tampering}) &= 0.85 \\ P(\text{alarm}|\neg\text{fire} \wedge \neg\text{tampering}) &= 0.0001 \\ P(\text{leaving}|\text{alarm}) &= 0.88 \\ P(\text{leaving}|\neg\text{alarm}) &= 0.001 \\ P(\text{report}|\text{leaving}) &= 0.75 \\ P(\text{report}|\neg\text{leaving}) &= 0.01 \end{aligned}$$

The following is a probabilistic Horn abduction representation of this Bayesian network:

$$\begin{aligned} &\text{disjoint}([\text{fire}(\text{yes}) : 0.01, \text{fire}(\text{no}) : 0.99]). \\ &\text{smoke}(Sm) \leftarrow \text{fire}(Fi), c\_smoke(Sm, Fi). \\ &\text{disjoint}([c\_smoke(\text{yes}, \text{yes}) : 0.9, c\_smoke(\text{no}, \text{yes}) : 0.1]). \\ &\text{disjoint}([c\_smoke(\text{yes}, \text{no}) : 0.01, c\_smoke(\text{no}, \text{no}) : 0.99]). \\ &\text{disjoint}([\text{tampering}(\text{yes}) : 0.02, \text{tampering}(\text{no}) : 0.98]). \end{aligned}$$

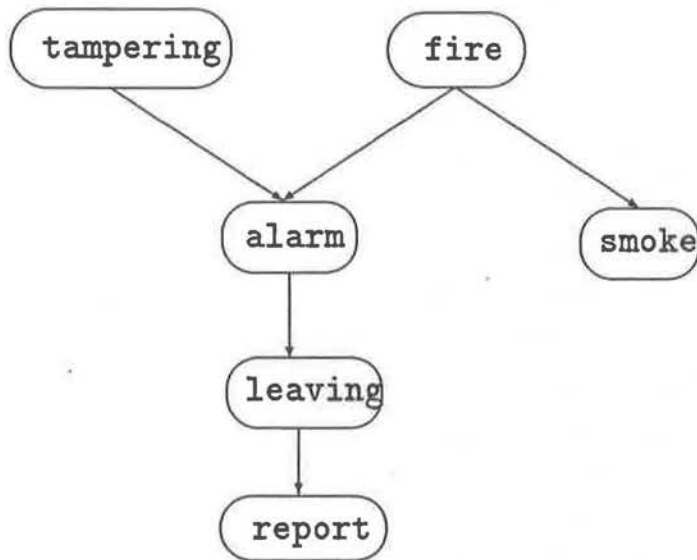


Figure 2: A Bayesian network for a smoking alarm.

$alarm(Al) \leftarrow fire(Fi), tampering(Ta), c\_alarm(Al, Fi, Ta).$   
 $disjoint([c\_alarm(yes, yes, yes) : 0.50, c\_alarm(no, yes, yes) : 0.50]).$   
 $disjoint([c\_alarm(yes, yes, no) : 0.99, c\_alarm(no, yes, no) : 0.01]).$   
 $disjoint([c\_alarm(yes, no, yes) : 0.85, c\_alarm(no, no, yes) : 0.15]).$   
 $disjoint([c\_alarm(yes, no, no) : 0.0001, c\_alarm(no, no, no) : 0.9999]).$   
 $leaving(Le) \leftarrow alarm(Al), c\_leaving(Le, Al).$   
 $disjoint([c\_leaving(yes, yes) : 0.88, c\_leaving(no, yes) : 0.12]).$   
 $disjoint([c\_leaving(yes, no) : 0.001, c\_leaving(no, no) : 0.999]).$   
 $report(Le) \leftarrow leaving(Al), c\_report(Le, Al).$   
 $disjoint([c\_report(yes, yes) : 0.75, c\_report(no, yes) : 0.25]).$   
 $disjoint([c\_report(yes, no) : 0.01, c\_report(no, no) : 0.99]).$

Note that here instead of creating  $c\_fire$  and making it equivalent to  $fire$ , we just made  $fire$  a hypothesis.

### 3.1 Equivalence Results

The basic equivalence result is the equivalence between joint distributions and explanations of the RVs having particular values:

Let  $P$  be the probability function sanctioned by a Bayesian network. Let  $T$  be the corresponding probabilistic Horn abduction theory. Let  $P_T$  be the probability function defined by the translation of a Bayesian network into the probabilistic Horn abduction framework. The aim is to show that these are the same.

The definition of the Bayes net distribution is using the joint probability for all values in the network. This corresponds to a conjunction of values for all RVs.

**Lemma 3.2** Suppose  $a_1, \dots, a_n$  are all of the RVs in a Bayesian network, with  $T$  as the corresponding probabilistic Horn abduction theory, then

$$P(a_1 = v_1 \wedge \dots \wedge a_n = v_n) = P_T(a_1(v_1) \wedge \dots \wedge a_n(v_n))$$

**Proof:** By definition of a Bayes net:

$$P(a_1 = v_1 \wedge \dots \wedge a_n = v_n) = \prod_{i=1}^n P(a_i = v_i | a_{i_1} = v_{i_1}, \dots, a_{i_{n_i}} = v_{i_{n_i}})$$

By definition of  $c_{.a_i}$  in the translation

$$P(a_i = v_i | a_{i_1} = v_{i_1}, \dots, a_{i_{n_i}} = v_{i_{n_i}}) = P_T(c_{.a_i}(v_i, v_{i_1}, \dots, v_{i_{n_i}}))$$

There is only one explanation of  $a_1(v_1) \wedge \dots \wedge a_n(v_n)$ , namely  $\{c_{.a_i}(v_i, v_{i_1}, \dots, v_{i_{n_i}}) : i = 1..n\}$ . Thus we have

$$\begin{aligned} P(a_1 = v_1 \wedge \dots \wedge a_n = v_n) &= \prod P(a_i = v_i | a_{i_1} = v_{i_1}, \dots, a_{i_{n_i}} = v_{i_{n_i}}) \\ &= \prod P_T(c_{.a_i}(v_i, v_{i_1}, \dots, v_{i_{n_i}})) \\ &= P_T(a_1(v_1) \wedge \dots \wedge a_n(v_n)) \end{aligned}$$

□

The equivalence between  $P$  and  $P_T$  now follows directly from theorem A.11. These two measures agree on the combinations of values for RVs, they both obey the probability axioms, and so they agree on all formulae.

**Theorem 3.3** If  $H$  is a set of assignments to random variables in a Bayesian Network, and  $H'$  is the analogous propositions to  $H$  in the corresponding probabilistic Horn abduction theory  $T$ , then

$$P(H) = P_T(H').$$

This should not be too surprising as the set of explanations of any formula correspond to an assignment of values for the ancestors in the Bayesian network. We can compute the values of any hypothesis by summing over the values of the ancestors of the hypothesis.

Appendix A gives a more formal development of this result.

### 3.2 Propositional Abduction in Bayesian Networks

The preceding section showed how any Bayesian network can be represented directly in terms of (propositional) probabilistic Horn abduction.

The opposite is also true. Every propositional probabilistic Horn abduction theory corresponds directly to a Bayesian network. Here we give the mapping.



Each disjoint declaration maps to a random variable. These form leaves (they have no ancestors) of the Bayesian network.

Every atom defined by rules also corresponds to a random variable.

If we have rules for  $a$

$$\begin{aligned} a &\leftarrow b_{11} \wedge \cdots b_{1n_1}. \\ &\vdots \\ a &\leftarrow b_{k1} \wedge \cdots b_{kn_k}. \end{aligned}$$

we can define atom  $r_i$  to correspond to the  $i$ -th body, and have rules  $a \leftarrow r_1$  through  $a \leftarrow r_k$  as well as  $k$  rules of the form  $r_i \leftarrow b_{i1} \wedge \cdots b_{in_i}$ .

We also create nodes for the  $r_i$  (we don't have to but it makes the contingency tables simpler).

We make arcs between going from the  $b_{ij}$  to  $r_i$ , and give the contingency table for a conjunction ("noisy-and"). We make arcs going from the  $r_i$  to  $a$ , giving the contingency table for a disjunction ("noisy-or").

When we do this mapping and then use the translation of the previous section to get back from the Bayesian network to a probabilistic Horn abduction theory, we get essentially the same probabilistic Horn abduction theory. This new theory has many new atoms (corresponding to negations, which may not have existed in the original theory, but must exist in the Bayesian network) which can be ignored. There also will be many more disjoint declarations, but these will correspond to extreme distributions (one value has probability 1), and can also be ignored (or partially evaluated away).

It is this direct correspondence that justifies our claim to having *the* logic of discrete Bayesian networks.

## 4 Discussion

### 4.1 Independence and dependence

It may seem at first that only allowing independent hypotheses places a restriction on what we can represent. People claim that there are dependencies amongst hypotheses in the world. The claim in this paper is that the world can be represented so that all of the hypotheses are independent. This apparent conflict can be resolved by noticing that the world does not determine

what the hypotheses are. Just as, when defining  $n$ -dimensional Euclidean spaces, we can define the space with non-orthogonal axes, we can also define the space with orthogonal axes. When we do so everything becomes simpler. Just because we can axiomatise a world using dependent hypotheses does not mean that we cannot define the world using independent hypotheses.

The justification for this claim is based on noticing that Bayesian networks (that can represent arbitrary probabilistic interaction) can be represented in our framework that uses only independent hypotheses. Note however that, as there can be an exponential number of independent values given  $n$  probabilistic hypotheses, we may have to create an exponential number of independent hypotheses in the probabilistic Horn abduction framework. We are not claiming that we are getting something for nothing.

Note that others have also noticed the universality of just having independent hypotheses. For example, consider Reichenbach's *principle of the common cause*:

“If coincidences of two events  $A$  and  $B$  occur more frequently than their independent occurrence, ... then there exists a common cause for these events ...” [49, p. 163].

When there is a dependency amongst random variables, we invent a hypothesis to explain that dependence. Thus the assumption of independence, while it gives a restriction on the knowledge bases that are legal, really gives no restriction on the domains that can be represented.

While we have the ability invent hypotheses, we don't need to consider non-independent hypotheses. I would argue that it is much simpler and more natural to invent new hypotheses to explain dependence rather than having to worry about dependence in the language.

## 4.2 Abduction and Prediction

When computing the prior probability of a hypothesis, we find the explanations for that hypothesis. This corresponds to the use of “abduction” [39].

If we consider conditional probability, as we normally do, we have

$$P(\alpha|\beta) = \frac{P(\alpha \wedge \beta)}{P(\beta)}$$

$$= \frac{\sum_{e_i \in \text{expl}(\alpha \wedge \beta, T)} P(e_i)}{\sum_{e_i \in \text{expl}(\beta, T)} P(e_i)}$$

We can generate  $\text{expl}(\alpha \wedge \beta, T)$  by explaining  $\alpha$  from the elements of  $\text{expl}(\beta, T)$ . This corresponds to the combination of abducing to causes and default reasoning to predictions from these causes [39, 41, 55]. Thus it should not be surprising that, as Pearl says

Poole [1987] has ... devised a logic-based system where default rules are restricted to the causal variety, and reasoning from evidence to hypotheses is accomplished by specialized "theory formation" procedures. Such causal-based systems enjoy the features of parsimony, stability and modularity and facilitate a more natural, declarative representation of world knowledge. [32, p. 502]

In the light of the results of this paper, this result may follow from Pearl's other arguments that Bayesian networks form the right characterisation of Causal reasoning. Abducing the causes and then assumption-based reasoning from causes to predicting what should follow, is the common feature of both Bayesian networks (see also, for example, Shachter and Heckerman [54]) and recent assumption-based logical schemes [39, 41, 55].

There is another close similarity between the abductive approaches and the network propagation scheme of Pearl [32]. Finding the explanations of some  $g$  corresponds to working "up" the Bayesian network from  $g$ . Given evidence  $\beta$ , we first find  $\text{expl}(\beta, T)$  (or at least some of the explanations [42, 44]). This involves searching up the tree from  $\beta$ . The next step involves finding  $\text{expl}(\alpha \wedge \beta, T)$ . This can be obtained from explaining  $\alpha$  from the explanations of  $\beta$ . If we want to compute this for all  $\alpha$ , we can do this by working down the tree from the explanations of  $\beta$ . This 2-phase approach is analogous to Pearl's network propagation scheme, with the initial moving up the tree corresponding to  $\lambda$  messages, and the second phase of moving down the tree from the explanations corresponds to the  $\pi$  messages of Pearl [32]. When the network is singly connected, we only need to send a summary of the explanations, and both phases can go in parallel, with a node able to send both  $\lambda$  and  $\pi$  messages at the same time. Given this analogy, it is also easy to see why the upward  $\lambda$  messages result in both  $\lambda$  and  $\pi$  messages (as we need to carry out both phases of the computation of the conditional

probability), while the  $\pi$  messages only result in other  $\pi$  messages (we are only in the second phase of computing the conditional probability, namely finding the explanations from those explanations of the observations).

### 4.3 Causation

There have been problems associated with logical formulations of causation [33]. There have been claims that Bayesian networks provide the right independencies for causation [34]. This paper provides evidence that abducing to causes and making assumptions as to what to predict from those assumptions [39, 41] is the right logical analogue of the independence in Bayesian networks (as described in section 4.2).

One of the problems in causal reasoning that Bayesian networks overcome [34] is in preventing reasoning such as “if  $c_1$  is a cause for  $a$  and  $c_2$  is a cause for  $\neg a$ , then from  $c_1$  we can infer  $\neg c_2$ ”. This is the problem that occurs, for example, in the Yale shooting problem [18]. Our embedding says that this does not occur in Bayesian networks as  $c_1$  and  $c_2$  must already be known to be disjoint. We must have already disambiguated what occurs when they are both true. If we represent the Yale shooting scenario so that the rules for “alive” are disjoint the problem does not arise.

$$\begin{aligned} \text{alive3}(\text{yes}) &\leftarrow \text{alive2}(\text{yes}) \wedge \text{shoot2}(\text{no}). \\ \text{alive3}(\text{yes}) &\leftarrow \text{alive2}(\text{yes}) \wedge \text{shoot2}(\text{yes}) \wedge \text{loaded2}(\text{no}). \\ \text{alive3}(\text{no}) &\leftarrow \text{alive2}(\text{no}). \\ \text{alive3}(\text{no}) &\leftarrow \text{alive2}(\text{yes}) \wedge \text{shoot2}(\text{yes}) \wedge \text{loaded2}(\text{yes}). \\ \text{loaded2}(\text{yes}) &\leftarrow \text{loaded1}(\text{yes}). \\ \text{loaded2}(\text{no}) &\leftarrow \text{loaded1}(\text{no}). \\ \text{alive2}(\text{yes}) &\leftarrow \text{alive1}(\text{yes}). \\ \text{alive2}(\text{no}) &\leftarrow \text{alive1}(\text{no}). \end{aligned}$$

It follows from this that given  $\text{alive1}(\text{yes})$ ,  $\text{loaded1}(\text{yes})$  and  $\text{shoot2}(\text{yes})$ , we can prove  $\text{alive3}(\text{no})$ .

I do not believe that this solution would have satisfied Hanks and McDermott [18], as we do not have general rules about persistence.

Probabilistic Horn abduction does not allow a general frame rule, even for a particular predicate such as

$$alive(T + 1) \leftarrow alive(T) \wedge alive\_persists(T)$$

since each defeater must be disjoint.

This formulation, however does get other persistence problems right. For example, where there is likelihood that the position of a car does not persist, because the car was stolen [21]. Thus it seems like a good foundation to look for general rules that allow us to build the contingency tables.

## 5 Representational Methodology

Once we have a tool, it is important to know how to use it. The problem of a representational methodology [41] is an important and much overlooked part of automated reasoning research.

It may seem that the assumptions used in designing probabilistic Horn abduction were so restrictive that the system would be useless for real problems. In this section, I argue that this is not the case.

The general idea is to use definite clauses to write a simulation (in the “causal” direction [54]) based on different possible hypotheses. This axiomatisation must follow the assumptions about the rule base and about the independence of hypotheses, but we argue in this section that this is not too difficult.

Indeed it is arguable, that rather than stifling the imagination of the axiomatiser to write “what is true” in their domain, placing restrictions on the representation language provides guidance to how to go about thinking about the domain. One of the aims of restricting the language is to make it easier to write and understand axiomatisations of the world. Whether this is true in practice, however, remains to be seen.

### 5.1 Disjoint and Covering Explanations

For our probabilistic analysis (section 2.4), we assumed that the explanations were disjoint and covering. If we want our probabilities to be correct, we must ensure that the rules for an atom are disjoint and covering.

If the rules for an atom  $a$  are not covering, we can invent another cause for the goal representing “all the other possible causes” of the atom [14, 40], and add

$$a \leftarrow a\_true\_for\_some\_other\_reason.$$

and make  $a\_true\_for\_some\_other\_reason$  into a hypothesis.

Although disjointedness of rules places a restriction on the knowledge base, it does not place a restriction on the sorts of knowledge that we can represent. In general, suppose we have rules<sup>4</sup>:

$$\begin{aligned} a &\leftarrow b_1. \\ &\vdots \\ a &\leftarrow b_n. \end{aligned}$$

Suppose we have the propositions  $\bar{b}_i$  to be the proposition that is the negation of  $b_i$  (see section 2.5.3), we can make sure the rules are disjoint by transforming them into

$$\begin{aligned} a &\leftarrow b_1. \\ a &\leftarrow \bar{b}_1 \wedge b_2. \\ a &\leftarrow \bar{b}_1 \wedge \bar{b}_2 \wedge b_3. \\ &\vdots \\ a &\leftarrow \bar{b}_1 \wedge \cdots \wedge \bar{b}_{n-1} \wedge b_n. \end{aligned}$$

Thus we make the rules disjoint, by ordering the rules and making sure that the bodies of rules are false if the bodies of previous rules are true.

Syntactically, this seems to increase the complexity of  $n$  rules to have  $\frac{n(n+1)}{2}$  atoms in the body. While this is true, there are only  $2n - 1$  different atoms that need to be explained. Thus, in practice, the complexity increases linearly not as a square.

**Example 5.1** Suppose we want to represent an “and-gate” that should have value 0 if either of the inputs are zero. Suppose we represent the proposition that port  $G$  has output  $V$  at time  $T$  as  $val(G, V, T)$ . We can ensure that the

<sup>4</sup>This corresponds to the “noisy or” [34, p. 184] of the  $b_i$ .

explanations are disjoint locally by ensuring that only one body can ever be true:

$$\begin{aligned}
 \text{val}(\text{out}(G), \text{off}, T) &\leftarrow \text{and\_gate}(G) \wedge \text{ok}(G) \\
 &\quad \wedge \text{val}(\text{input}(1, G), \text{off}, T). \\
 \text{val}(\text{out}(G), \text{off}, T) &\leftarrow \text{and\_gate}(G) \wedge \text{ok}(G) \\
 &\quad \wedge \text{val}(\text{input}(1, G), \text{on}, T) \\
 &\quad \wedge \text{val}(\text{input}(2, G), \text{off}, T). \\
 \text{val}(\text{out}(G), \text{on}, T) &\leftarrow \text{and\_gate}(G) \wedge \text{ok}(G) \\
 &\quad \wedge \text{val}(\text{input}(\text{on}, G), 1, T) \\
 &\quad \wedge \text{val}(\text{input}(2, G), \text{on}, T).
 \end{aligned}$$

Note that the third conjunct in body of the second rule (the  $\text{val}(\text{input}(1, G), \text{on}, T)$ ) is there to ensure that the bodies are disjoint.

This has repercussions in biasing the most likely explanation to the first rule, which is more general than the others. To make it more fair the first rule could be split into two cases depending on the value of input 2. This problem of the most likely diagnosis depending on the representation seems endemic to approaches that try to find the diagnosis (either explanation or interpretation) that is “most likely” [34, 46].

## 5.2 Causation Events

When representing knowledge for abduction [40, 41], we have to be able to make sure that we can imply the observations. In general a fault or disease doesn't *imply* a particular observation. For example, having a cold does not imply sneezing, but could cause sneezing. A gate being in an unknown state does not imply any particular value for the output of the gate. To solve this problem we introduce another hypothesis that the cold caused the sneezing. In the other example, we have to hypothesise that the gate is producing a particular value. This idea is analogous to the notion of a “causation event” of Peng and Reggia [35].

The cold causing sneezing could be written as

$$\text{sneeze} \leftarrow \text{cold} \wedge \text{cold\_caused\_sneeze}$$

Following Peng and Reggia [35], one way to implement the causation events, is to use the relations *has\_disease(D)* to mean that the patient has disease *D*; *actually\_causes(D, M)* to mean that disease *D* “actually caused” manifestation *M*; and *has\_manifestation(M)* to mean that the patient has manifestation *M*.

We can say that a manifestation is caused by the disease that actually causes it by:

$$\begin{aligned} \text{has\_manifestation}(M) \leftarrow & \text{has\_disease}(D) \\ & \wedge \text{actually\_causes}(D, M). \end{aligned}$$

The conjunction

$$\text{has\_disease}(D) \wedge \text{actually\_causes}(D, M)$$

corresponds to Peng and Reggia’s [35] causation event  $M : D$ .

We have the disjoint declarations for each  $i, j$ :

$$\text{disjoint}([\text{actually\_causes}(d_i, m_j) : p_{ij}, \text{didnt\_actually\_cause}(d_i, m_j) : q_{ij}])$$

where  $p_{ij}$  corresponds to the the “conditional causal probability” (“causal strength”) of [35], and  $q_{ij} = 1 - p_{ij}$ .  $p_{ij}$  can be seen as the fraction of the cases where  $d_i$  is true that  $d_i$  “actually causes”  $m_j$ .

We also have the possible hypotheses

$$\text{disjoint}([\text{has\_disease}(d_i) : p_i, \text{doesnt\_have\_disease}(d_i) : q_i])$$

where  $p_i$  is the prior probability of the disease  $d_i$ , and  $q_i = 1 - p_i$ .

To implement this we still have to worry about making the rules disjoint. This is done in the same way as in section 5.1. If manifestation  $m$  has possible causes  $d_1, \dots, d_k$ , we write:

$$\begin{aligned} \text{has\_manifestation}(m) \leftarrow & \text{has\_disease}(d_1) \\ & \wedge \text{actually\_causes}(d_1, m). \\ \text{has\_manifestation}(m) \leftarrow & \text{has\_disease}(d_2) \wedge \text{doesnt\_have\_disease}(d_1) \\ & \wedge \text{actually\_causes}(d_2, m). \\ & \vdots \end{aligned}$$



$$\begin{aligned}
has\_manifestation(m) \leftarrow & has\_disease(d_k) \\
& \wedge doesnt\_have\_disease(d_1) \wedge \dots \\
& \wedge doesnt\_have\_disease(d_{k-1}) \\
& \wedge actually\_causes(d_k, m).
\end{aligned}$$

The advantage of making these disjoint is that we are now able to interpret “actually causing” observationally for the cases where there are two possible causes. Here we arbitrarily assign the “actual cause” to the first disease. It is now easy to interpret the notion of “actually causing”, as there is no ambiguity in any data. This makes the concept of “actual cause” into an observational notion for which we can collect statistics and do not need a theory of causation that is deeper than the theory we want to represent.

### 5.2.1 Hypotheses with indeterminate output

There is one case where we have to be concerned about causation events as well as the problem of parametrizing possible hypotheses and the interaction with the independence assumption. I have argued elsewhere [40, 41] that there is much power obtainable and subtlety involved in parametrizing hypotheses appropriately. In this section we expand on previous analysis [41], and show how probabilities affect parametrization considerations when using causation events by considering some case studies.

As an example, suppose we have a gate  $G$  that takes two values as input, and outputs a value that can be in the range 1 to  $n$ . Suppose we want to represent the gate being in an unknown state (this is applicable whether or not we have fault models [14, 40]). Suppose we represent the proposition that gate  $G$  has output  $V$  at time  $T$  as  $val(G, V, T)$ .

We cannot represent the hypothesis that the gate is in the unknown state by using the hypothesis  $u(G)$  and the fact

$$val(out(G), V, T) \leftarrow u(G).$$

The problem is that the above fact states that a gate in the unknown state produces *all* values of output, rather than saying that it produces some output. Knowing a gate is in an unknown state does not imply any value for the output.

When there are no probabilities involved [41, 40] we parametrize the hypothesis by the values on which it depends. This could be done by having the

hypothesis  $produces(G, V, T)$  (interpreted as “gate  $G$  is faulty and produces value  $V$  at time  $T$ ”) and the rule

$$val(out(G), V, T) \leftarrow produces(G, V, T).$$

We would say that a port has only one value at a time by having the disjoint declaration<sup>5</sup>:

$$disjoint([ok(P) : p_0, produces(P, v_1, T) : p_1, \dots, produces(P, v_n, T) : p_n]).$$

Suppose we know that gate  $g_1$  has probability  $\epsilon$  of being in the unknown state. Then  $p_0 = 1 - \epsilon$ . If we assume that each possible output value has equal chance, and that there are  $n$  possible output values, then  $p_i$ , the prior probability that it produces output value  $v_i$  is  $\epsilon/n$  for  $1 \leq i \leq n$ .

When we have more than one observation, there is another problem. For the probabilities we assumed that the hypotheses were independent. We would not expect that

$$P(produces(g_1, 1, t_2) | produces(g_1, 1, t_1)) = P(produces(g_1, 1, t_2))$$

Once we know that the gate is in an unknown state at time  $t_1$  it should not be so unlikely that it is in an unknown state at time  $t_2$ . Put another way, once we have paid the price once for assuming that the gate is in an unknown state at time  $t_1$  we should not pay the price again for assuming that it is in an unknown state at time  $t_2$ .

To work in general, we need a mixture of the above two ideas. Suppose a gate  $G$  has probability of  $\epsilon$  of being in the unknown state, and that there are  $n$  possible output values, each of which has an equal prior chance of being produced by a gate in the unknown state. This can be represented as the hypotheses<sup>6</sup>

$$disjoint([produces(P, v_1, T) : \frac{1}{n}, \dots, produces(P, v_n, T) : \frac{1}{n}]).$$

$$disjoint([ok(G) : p_0, \dots, u(G) : \epsilon]).$$

<sup>5</sup>Here we have assumed that there are no fault states other than the unknown state. These could be added to this declaration without changing the point of the discussion.

<sup>6</sup>Here we assume a uniform distribution of values. Any other distribution could be given.

and the rule

$$val(out(G), V, T) \leftarrow u(G) \wedge produces(G, V, T).$$

$u(G)$  means  $G$  is in the unknown state, and  $produces(G, V, T)$  means that given gate  $G$  is broken, it produces value  $V$  at time  $T$ . The system assumes once that the gate is broken, and then makes other assumptions of what values it is producing at different times.

The atom  $produces(G, V, T)$  can be seen as a “causation event”, that we invent because being in an unknown fault state does not imply any particular value.

### 5.2.2 Intermittent versus non-intermittent faults

Because of the way we parametrized the hypotheses, the above representation of faults says that the output is only a function of the time. The hypothesis  $produces(G, V, T)$  and the above rules places no constraints on the values of the outputs at different times. This is a way to represent the fact that the gate can have an intermittent fault (it depends only on the time of observation). There is no constraint that says the gate produces the same output when given the same inputs at different times.

We can give the non-intermittency assumption by saying that the fault only depends on the input and not on the time. This can be done instead by having the hypothesis  $prod(G, V, I_1, I_2)$  (meaning gate  $G$  produces output  $V$  when given  $I_1$  and  $I_2$  as input) and a rule

$$\begin{aligned} val(out(G), V, T) \leftarrow & u(G) \wedge prod(G, V, I_1, I_2) \\ & \wedge val(input(1, G), I_1, T) \\ & \wedge val(input(2, G), I_2, T). \end{aligned}$$

## 5.3 Two Examples

In this section we show the complete theories for two non-trivial examples.

**Example 5.2** This first example is an implementation of cascaded one bit adders (Figures 3 and 4), to form a ripple adder. The axiomatisation is adapted from the consistency-based axiomatisation of Genesereth [16].

$val(P, V, T)$  means that port  $P$  has value  $V$  at time  $T$ . We use as simple a representation of time as is needed. In this case we need to be able to have different observations at different times, and use constants to denote different times.

We first axiomatise how gates work. We must axiomatise how normal gates as well as faulty gates work [38, 40].

Each of the gates can be in one of four states (ok, stuck on, stuck off or unknown). When the gate is in the unknown ( $ab$ ) state it can produce either value with equal probability. We also used the intermittency assumption.

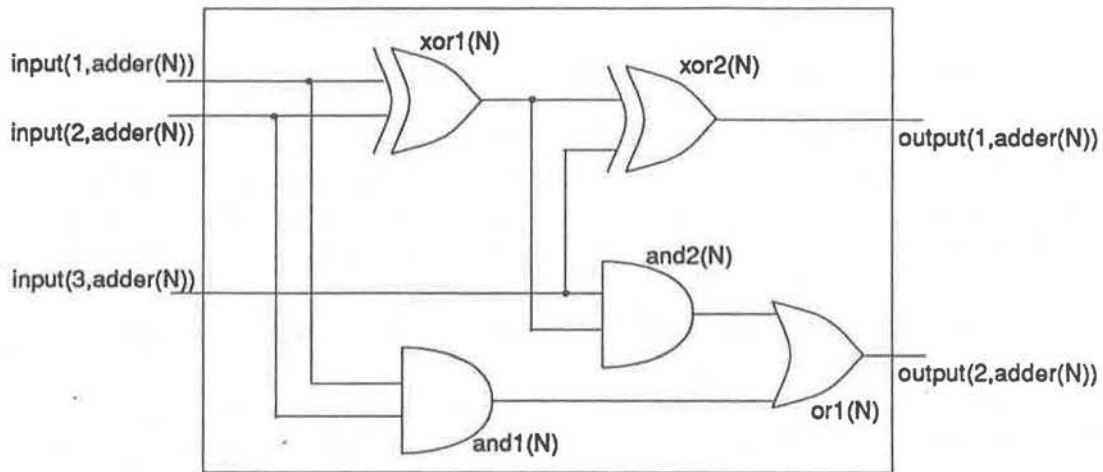


Figure 3: One bit adder,  $adder(N)$ .

```

val(output(G), off, T) ←
    gate(G, and), ok(G), val(input(1, G), off, T).
val(output(G), off, T) ←
    gate(G, and), ok(G), val(input(1, G), on, T),
    val(input(2, G), off, T).
val(output(G), on, T) ←
    gate(G, and), ok(G), val(input(1, G), on, T),
    val(input(2, G), on, T).
    
```




Figure 4: Three cascaded adders

```

val(output(G), on, T) ←
    gate(G, or), ok(G), val(input(1, G), on, T).
val(output(G), off, T) ←
    gate(G, or), ok(G), val(input(1, G), off, T),
    val(input(2, G), off, T).
val(output(G), on, T) ←
    gate(G, or), ok(G), val(input(1, G), off, T),
    val(input(2, G), on, T).

val(output(G), off, T) ←
    gate(G, xor), ok(G), val(input(1, G), off, T),
    val(input(2, G), off, T).
val(output(G), on, T) ←
    gate(G, xor), ok(G), val(input(1, G), off, T),
    val(input(2, G), on, T).
val(output(G), on, T) ←
    gate(G, xor), ok(G), val(input(1, G), on, T),
    val(input(2, G), off, T).
val(output(G), off, T) ←
    gate(G, xor), ok(G), val(input(1, G), on, T),
    val(input(2, G), on, T).

val(output(G), on, T) ← stuck1(G).
val(output(G), off, T) ← stuck0(G).
val(output(G), V, T) ← ab(G), produced(G, V, T).
    
```

$val(P, V, T) \leftarrow conn(Q, P), val(Q, V, T).$

Note how we have made the rules disjoint, by adding extra conditions. For example, the second rule includes the condition  $val(input(1, G), on, T)$  that is there just to make sure the rules are disjoint. Treated as a definite clause in isolation this rule is true without this second condition.

We also specify the random variables as outlines in the preceding section. Note that this implies that the gates fail independently.

$disjoint([ok(X) : 0.999, ab(X) : 0.0000001, stuck1(X) : 0.0004999, stuck0(X) : 0.0005]).$   
 $disjoint([produced(X, on, T) : 0.5, produced(X, off, T) : 0.5]).$

We axiomatise how the gates in an adder are connected, and what gates there are in an adder.

$conn(input(1, adder(N)), input(1, xor1(N))).$   
 $conn(input(1, adder(N)), input(1, and1(N))).$   
 $conn(input(2, adder(N)), input(2, xor1(N))).$   
 $conn(input(2, adder(N)), input(2, and1(N))).$   
 $conn(input(3, adder(N)), input(2, xor2(N))).$   
 $conn(input(3, adder(N)), input(1, and2(N))).$   
 $conn(output(xor1(N)), input(1, xor2(N))).$   
 $conn(output(xor1(N)), input(2, and2(N))).$   
 $conn(output(and1(N)), input(2, or1(N))).$   
 $conn(output(and2(N)), input(1, or1(N))).$   
 $conn(output(xor2(N)), output(1, adder(N))).$   
 $conn(output(or1(N)), output(2, adder(N))).$

$conn(output(2, adder(N)), input(3, adder(N1))) \leftarrow succ(N, N1).$   
 $val(input(3, adder(1)), off, T).$

$gate(xor1(N), xor).$   
 $gate(xor2(N), xor).$   
 $gate(and1(N), and).$   
 $gate(and2(N), and).$

$gate(or1(N), or)$ .

$succ(N, s(N))$ .

The relation  $succ(N, N1)$  is used to state when one gate is next to another. This allows us to observe arbitrarily large cascaded adders.

In order for us to be able to observe inputs and to be able to predict expected values from unknown inputs we can make the inputs to the gates to be random variables. (The alternative is to write as facts what the inputs to the gates are [40]).

$disjoint([val(input(1, adder(N)), on, T) : 0.5, val(input(1, adder(N)), off, T) : 0.5]).$   
 $disjoint([val(input(2, adder(N)), on, T) : 0.5, val(input(2, adder(N)), off, T) : 0.5]).$

We can specify an observation such as that  $10 + 11$  gave 001, as

$$\begin{aligned} & val(input(1, adder(1)), off, t1) \\ \wedge & val(input(1, adder(s(1))), on, t1) \\ \wedge & val(input(2, adder(1)), on, t1) \\ \wedge & val(input(2, adder(s(1))), on, t1) \\ \wedge & val(output(1, adder(1)), on, t1), \\ \wedge & val(output(1, adder(s(1))), off, t1) \\ \wedge & val(output(2, adder(s(1))), off, t1) \end{aligned}$$

**Example 5.3** The second example is of the framework for depiction and image interpretation of Reiter and Mackworth [53]. Here we interpret simple line drawings of a map. These consist of lines and areas that depict roads, rivers, shores, lakes and land. This axiomatisation is based on the abductive representation of [41]<sup>7</sup>.

We axiomatise how scene objects could have produced image objects. Given an image we conjecture scene elements that could have produced that image.

The main difference between this axiomatisation and those of [53, 41] is that we have to make constructive derivations of the image. Rather than

<sup>7</sup>As in [41], for simplicity, the depicts relation of [53] has not been used even though it facilitates the representation of some, more complicated, examples.

starting with all interpretations, and use consistency to prune those that are impossible, we make sure that we can only generate possible explanations. This follows the methodology given earlier in this section.

We first allow one to write the building blocks of explanations.  $area(X, land)$  means that the object  $X$  in the image depicts land in scene, given that it is a region. The probabilities reflect that in our (made-up) domain, 70% of the areas are water and 30% are land.

$region(X) \leftarrow area(X, T).$   
 $chain(X) \leftarrow linear(X, T).$   
 $disjoint([area(X, land) : 0.3, area(X, water) : 0.7]).$   
 $disjoint([linear(X, road) : 0.2, linear(X, river) : 0.5, linear(X, shore) : 0.3]).$

We now have axioms that describe how image objects could be produced in terms of scene objects.

$tee(X, Y, E)$  means that end  $E$  of chain  $X$  ends at chain  $Y$ . This can either be because  $X$  is a road that joins  $Y$ , or  $X$  is a river that starts at  $Y$ , or  $X$  is a river that flows into (river or shore)  $Y$ . We arbitrarily number ends of chains with a 0 or a 1, and one end of a river needs to be a mouth and one a source of the river.

$tee(X, Y, E) \leftarrow joins(X, Y, E) \wedge linear(X, road).$   
 $tee(X, Y, E) \leftarrow joins(X, Y, E) \wedge linear(X, river) \wedge$   
 $linear(Y, road) \wedge source(X, E).$   
 $tee(X, Y, E) \leftarrow linear(X, river) \wedge canflowto(Y) \wedge$   
 $flowsto(X, Y) \wedge mouth(X, E).$   
 $canflowto(X) \leftarrow linear(X, river).$   
 $canflowto(X) \leftarrow linear(X, shore).$   
 $disjoint([joins(X, Y, E) : 0.05, notjoins(X, Y, E) : 0.95]).$   
 $disjoint([mouth(X, 0) : 0.5, mouth(X, 1) : 0.5]).$   
 $disjoint([flowsto(X, Y) : 0.1, notflowsto(X, Y) : 0.9]).$   
 $disjoint([source(R, 1) : 0.5, source(R, 0) : 0.5]).$

Similarly we can handle two chains crossing (a chi —  $\chi$ ) in the image. Here again we just use the notion of a causal event to make the implication always true.

Figure 5 gives the correspondences between image and scene predicates.



scene	image
<i>linear(X, road)</i> <i>linear(X, river)</i> <i>linear(X, shore)</i>	<i>chain(X)</i>
<i>area(X, land)</i> <i>area(X, water)</i>	<i>chain(X)</i>
<i>joins(X, Y, E)</i> <i>flowsto(X, Y)</i>	<i>tee(X, Y, E)</i>
<i>docross(X, Y)</i>	<i>chi(X, Y)</i>
<i>source(X, N)</i> <i>petersout(X, N)</i>	<i>open(X, N)</i>
<i>linear(X, shore)</i> <i>roadloop(X)</i>	<i>closed(X)</i>
<i>beside(X, Y)</i>	<i>bounds(X, Y)</i>
<i>inside(X, Y)</i>	<i>interior(X, Y)</i>
<i>outside(X, Y)</i>	<i>exterior(X, Y)</i>

Figure 5: Image–scene predicates.

$chi(X, Y) \leftarrow crossable(X, Y) \wedge docross(X, Y).$   
 $crossable(X, Y) \leftarrow linear(X, XT) \wedge linear(Y, YT) \wedge crosstype(XT, YT).$   
 $crosstype(road, road).$   
 $crosstype(road, river).$   
 $crosstype(river, road).$   
 $crosstype(road, shore).$   
 $crosstype(shore, road).$   
 $disjoint([docross(X, Y) : 0.2, dontcross(X, Y) : 0.8]).$

We can also have a chain being open or closed. Note here that the probabilities tell us that all shores form closed loops, and some, but very few roads form loops. It is also more likely that a river ends nowhere than a road does.

$open(X, N) \leftarrow linear(X, river) \wedge source(X, N).$   
 $open(X, N) \leftarrow linear(X, road) \wedge petersout(X, N).$   
 $disjoint([petersout(X, E) : 0.1, doesntpeterout(X, E) : 0.9]).$   
 $closed(X) \leftarrow linear(X, shore).$   
 $closed(X) \leftarrow linear(X, road) \wedge roadloop(X).$   
 $disjoint([roadloop(X) : 0.01, notloop(X) : 0.99]).$

We can also have a chain bounding an area, and reason about what areas can be inside or outside loops. Note that the rule implying *exterior* assumes that *inside* is functional. In other words there is only one area (directly) inside a closed loop. If this is not the case another representation needs to be chosen.

$bounds(X, Y) \leftarrow linear(X, XT) \wedge area(Y, YT) \wedge$   
 $\quad beside(X, Y) \wedge possbeside(XT, YT).$   
 $possbeside(road, land).$   
 $possbeside(river, land).$   
 $possbeside(shore, land).$   
 $possbeside(shore, water).$   
 $disjoint([beside(X, Y) : 0.1, notbeside(X, Y) : 0.9]).$   
 $disjoint([inside(X, Y) : 0.1, outside(X, Y) : 0.1, noside(X, Y) : 0.8]).$   
 $interior(X, Y) \leftarrow inside(X, Y).$   
 $exterior(X, Z) \leftarrow outside(X, Z) \wedge inside(X, Y) \wedge linear(X, XT),$   
 $\quad area(Y, YT) \wedge area(Z, ZT) \wedge possreg(YT, XT, ZT).$   
 $possreg(land, road, land).$

*possreg(land, shore, water).*  
*possreg(water, shore, land).*

An image becomes an observation that we condition on. For example, the image of figure 6 is represented as the observation

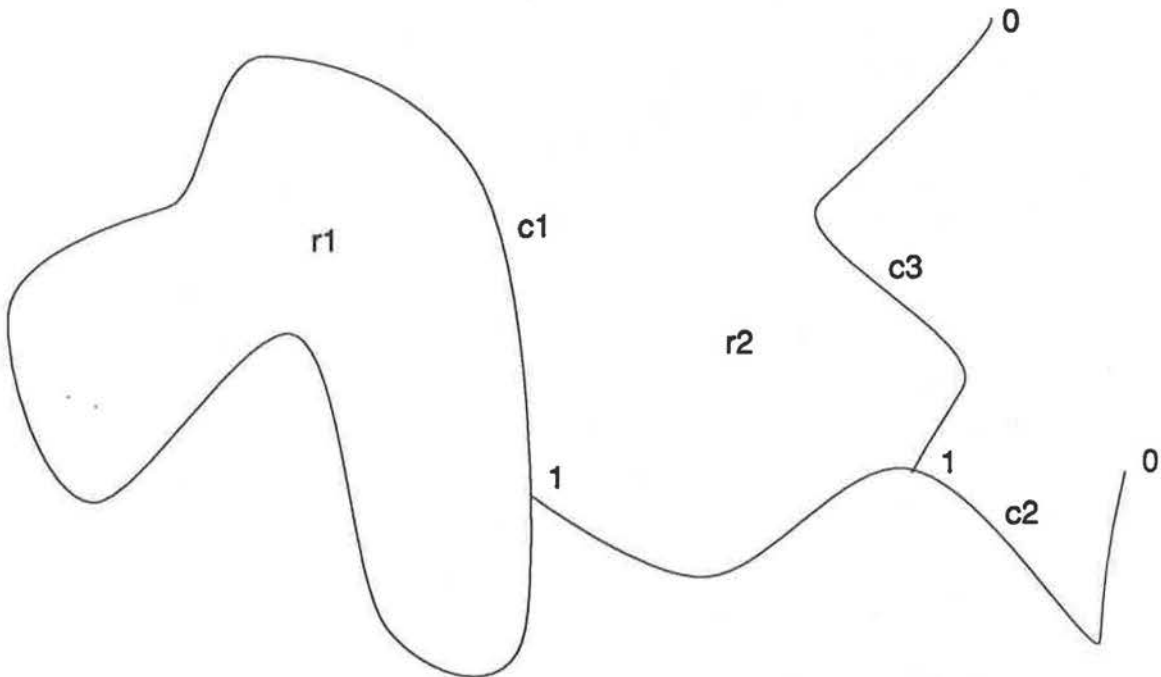
$$\begin{aligned} &chain(c1) \wedge chain(c2) \wedge chain(c3) \wedge region(r1) \wedge region(r2) \wedge \\ &tee(c2, c1, 1) \wedge bounds(c2, r2) \wedge bounds(c1, r1) \wedge bounds(c1, r2) \\ &\wedge exterior(c1, r1) \wedge interior(c1, r2), open(c2, 0) \wedge closed(c1) \wedge \\ &open(c3, 0) \wedge tee(c3, c2, 1) \wedge bounds(c3, r2) \end{aligned}$$


Figure 6: A simple image.

#### 5.4 Arbitrary Individuals

One of the problems that is considered in [20], is that of when there can be arbitrary individuals that affect a value, and the individuals present can only

be determined at run time.

As an example, consider the problem of a fire alarm going off, where it goes off if it was set off by one of the individuals present. All of the individuals can independently set off the alarm. We cannot write rules such as

$$\text{alarm}(\text{sounds}) \leftarrow \text{present}(P) \wedge \text{set\_off\_alarm}(P).$$

as the disjoint rules condition is violated. This is related to the problem in first order logic of being able to count the number of people present given just a database of relations of the form  $\text{present}(P)$ . We cannot prove there are 5 people present just because we have  $\text{present}(P)$  true for 5 instances of  $P$ , unless we have stated that everyone else is not present. Like pure Prolog, we cannot represent such knowledge.

We can however encode the problem so that we can count the number of people present. This is by forcing us to write  $\text{present}(L)$  where  $L$  is a list of the people present.

We can now represent the problem by<sup>8</sup>:

$$\begin{aligned} \text{alarm}(\text{sounds}) &\leftarrow \text{present}(L) \wedge \text{one\_set\_off\_alarm}(L). \\ \text{alarm}(\text{quiet}) &\leftarrow \text{present}(L) \wedge \text{none\_set\_off\_alarm}(L). \\ \text{one\_set\_off\_alarm}([H|T]) &\leftarrow \text{set\_off\_alarm}(H). \\ \text{one\_set\_off\_alarm}([H|T]) &\leftarrow \text{didnt\_set\_off\_alarm}(H) \\ &\quad \wedge \text{one\_set\_off\_alarm}(T). \\ \text{none\_set\_off\_alarm}([H|T]) &\leftarrow \text{didnt\_set\_off\_alarm}(H) \\ &\quad \wedge \text{none\_set\_off\_alarm}(T). \\ \text{none\_set\_off\_alarm}([]). \end{aligned}$$

$$\text{disjoint}([\text{set\_off\_alarm}(P) : p_1, \text{didnt\_set\_off\_alarm}(P) : p_2]).$$

Here  $p_1$  is the probability that a person would have set off an alarm given that no one before them had already set off the alarm.  $p_2 = 1 - p_1$ .

<sup>8</sup>Here we use the Prolog syntactic sugar for lists.  $[]$  is just a binary function symbol,  $[]$  is a constant, and the notational convention that  $[\alpha|[\beta]]$  is written as  $[\alpha, \beta]$  for any sequence of symbols  $\alpha$  and  $\beta$ .

## 6 Comparison with Other Systems

### 6.1 Other logic-based abductive schemes

There are many other proposals for logic-based abduction schemes (e.g., [47, 10, 17]). These however consider that we have to find all of the diagnoses. In practice there are prohibitively many of these. It is also not clear what to do with all of the explanations; there are too many to give to a user, and the costs of tests to determine which of the diagnoses is the “real” diagnosis is usually not outweighed by the advantages of finding the real diagnosis (see Ledley and Lusted [25] for an early description of the importance of probabilistic and value information in diagnosis). We provide an answer to both of these, namely using the explanations to compute probabilities, and not generating unlikely explanations [44].

The closest version of abduction to that presented here is that of Goebel et. al. [17], where there is also the simple abductive scheme where we do not need to do any chaining in order to determine inconsistency. Essentially we have added probabilities to that scheme under certain assumptions about the knowledge base and independence.

### 6.2 Probability and diagnosis

de Kleer and Williams [13, 14] and Peng and Reggia [35] both incorporate probabilistic knowledge to find the most likely diagnoses, but do not provide as flexible and simple a representation language as the one here.

de Kleer and Williams [13, 14] have explored the idea of using probabilistic information in consistency-based diagnosis (see [38, 40] for comparisons between abductive and consistency-based diagnoses).

The major differences between their approach and the one presented in this paper is that they differ in what they compute the probability of. de Kleer and Williams are finding the most likely interpretations (assignment of values to all hypotheses). This is the same as the diagnoses of Peng and Reggia [35] and the composite beliefs of Pearl [32], but is different to the diagnoses of de Kleer, Mackworth and Reiter [12]. We are computing the most likely explanations; we want to remain agnostic about the value of the irrelevant hypotheses. Both schemes allow us to ignore the unlikely diagnoses. de Kleer and Williams cannot distinguish between those remaining diagnoses

that differ in substantial ways from the most likely interpretations, and those that differ only in varying values that are irrelevant to the diagnosis. In our system, hypotheses that are not part of an explanation are irrelevant and are ignored. We do not place such an importance on the explanations, but rather on using the explanations to compute probabilities.

Peng and Reggia [35] also consider an abductive definition of diagnosis and incorporate probabilities, and best-first search. One difference is that they are trying to find probabilities of interpretations, but we are using explanations to find the probabilities of atoms. The main difference is in the underlying language. They use the notion of "hyper-bipartite" graphs made up of causation relations on sets of manifestations (can be observed), disorders (can be hypothesised), and pathological states.

One way to look at what they are doing is to consider it as a restriction of the system presented here where the language is propositional and only allows one element in the body of a clause. It is, however, not expressed in a logical language.

Hobbs et. al. [19] have devised a "cost-based abduction" for interpretation of natural language. Their scheme is similar to the one presented here, but they use costs associated with assumptions rather than probabilities. These costs can be seen as  $-\log$  probabilities [7]. One can view the current work as extending Hobbs et. al.'s to derive posterior probabilities in a consistent manner.

### 6.3 Logic and Bayes nets

The representation of Bayes nets is related to the work by Charniak and Shimony [7, 56]. Instead of considering abduction, they consider models that consist of an assignment of values to each random variable. The *label* of [56] plays an analogous role to our hypotheses. They however, do not use their system for computing posterior probabilities. It is also not so obvious how to extend their formalism to more powerful logics.

Horsch and Poole [20], Breese [5] have defined systems that incorporate Prolog style rules and Bayesian networks. These were designed to allow for dynamic construction of Bayesian networks. The rules of [20] cannot be interpreted logically but are macros that map into Bayesian network structure. The rules of [5] are at the meta-level of the Bayesian network, and build the network. We treat the rules as the object-level knowledge representation. It

is not clear how these proposals shed light on the underlying structure of the Bayesian network; they could just as easily be used to build any other propositional structure.

#### 6.4 Horn abduction and Dempster-Shafer

This work is also closely related to recent embeddings of Dempster-Shafer theory in ATMS [24, 48]. One difference between our embedding of Bayesian networks and Dempster-Shafer is in the independence assumptions used. Dempster-Shafer theory assumes that different rules are independent. We assume they are exclusive. Another difference is that these embeddings do not do evidential reasoning (by doing abduction), determining probability of hypotheses given evidence, but rather only determine the "belief" of propositions from forward chaining.

#### 6.5 Argument systems

Doyle [15], Loui [29] have argued that decisions can be best seen in terms of arguments for and against some propositions. Other have viewed nonmonotonic reasoning in terms of arguments [37, 26, 28, 36]. The explanations that we use can be seen as premises for logical arguments. We can determine the probability of some hypothesis by coming up with arguments for the proposition. We may only need some arguments to give an estimate of the probability of the hypothesis (see [42, 44]).

### 7 Conclusion

This paper has presented a pragmatically-motivated simple logic formulation that includes definite clauses and probabilities over hypotheses. This was designed to be a compromise between representational adequacy, ease to interpret semantically what the knowledge means, and ease of implementation [44]. It is suggested that this simple tool provides a good representation for many evidential reasoning tasks.

This is supported by the demonstration that probabilistic Horn abduction arguably forms *the* logic of discrete Bayesian networks. There is a direct

mapping between the knowledge in a Bayes net and the knowledge in a probabilistic Horn abduction theory.

This is also interesting because it provides a link to earlier work on the use of assumption-based reasoning for default reasoning [37, 39, 41]. One of the ways of viewing default reasoning is where an adversary chooses the assumptions. One way of viewing the probabilistic Horn abduction is as an instance of assumption-based reasoning, but where nature chooses the assumptions. This paper also demonstrates the correspondence between the observations that need to be explained in abduction [40], and what is conditioned on in Bayesian probability.

## A Formal Semantics

In this section we give the formal semantics for our language. As the language is very simple, the semantics will be correspondingly simple. The semantics will basically be that of Bacchus [3], restricted to our language, and incorporating our assumptions.

The logical statements will restrict the possible worlds, and the probabilities will provide measures over the possible worlds. The language is specially designed so that the logical facts neither implies a hypothesis nor implies the negation of a hypothesis. This means that we can treat the logic part and the probabilistic part of our semantics independently.

The only tricky thing we have to worry about is that there are potentially infinitely many independent hypotheses, with non-extreme (i.e., not equal to 0 or 1) probabilities. If we have one function symbol and one parametrized hypothesis  $h(X)$  with non extreme probability, we have infinitely many independent hypotheses of the form  $h(t)$  for each ground term  $t$ . Once we have this, the probability of each possible world will be zero. Thus, we cannot just sum over the possible worlds to determine the probability of a proposition (as is done in [3]). We instead provide a measure over sentences that can be described in our language (as in [2]). As we have made sure that the logic provides no constraints on the probabilities, we only need to consider sentences made of hypotheses.

As discussed in section 2.4, we make the unique names assumption. This is specified formally by making the domain we consider be the Herbrand Universe [27].



**Definition A.1** A semantic structure is a tuple  $\langle W, D, \phi, \pi, P^* \rangle$ , where  $W$  is a non-empty set. Elements of  $W$  are called possible worlds.

$D$  is a non-empty set (of individuals).  $D$  here is the Herbrand Universe (the set of ground terms in the language).

$\phi$  is a function that maps each  $n$ -ary function symbol to an element of  $D^n \mapsto D$  (in particular  $\phi$  maps each constant to an element of  $D$ ). We can extend  $\phi$  to ground terms, by the use of the recursive scheme  $\phi(f(t_1, \dots, t_n)) = \phi(f)(\phi(t_1), \dots, \phi(t_n))$ . In particular  $\phi$  is the identity function so that  $\phi(t) = t$  for any term  $t$ . Because the mapping does not depend on the world, these form "rigid designators".

$\pi$  is a 1-1 function such that  $\pi(w)$  maps each  $n$ -ary predicate symbol into a subset of  $D^n$ .

$P^*$  is a function from  $H \mapsto [0, 1]$ .

**Definition A.2** We define the semantic relation  $\models_{\langle W, D, \phi, \pi, P^* \rangle}$  between possible worlds and formulae:

The first case is for atoms

$$w \models_{\langle W, D, \phi, \pi, P^* \rangle} p(t_1, \dots, t_n) \text{ if } \langle \phi(t_1), \dots, \phi(t_n) \rangle \in \pi(w)(p).$$

The second is for conjunctions between atoms in the bodies of rules, as well as for conjunctions of definitions in our theory:

$$w \models_{\langle W, D, \phi, \pi, P^* \rangle} b_1 \wedge b_2 \text{ iff } w \models_{\langle W, D, \phi, \pi, P^* \rangle} b_1 \text{ and } w \models_{\langle W, D, \phi, \pi, P^* \rangle} b_2$$

The third rule defines the truth of definitions, and incorporates the assumptions about disjointedness and coveringness:

If  $b_1, \dots, b_n$  are all the bodies defining  $a$

$$w \models_{\langle W, D, \phi, \pi, P^* \rangle} (a \leftarrow b_1) \wedge \dots \wedge (a \leftarrow b_n)$$

$$\text{if } (w \models_{\langle W, D, \phi, \pi, P^* \rangle} a \text{ and } \exists i (w \models_{\langle W, D, \phi, \pi, P^* \rangle} b_i) \text{ and } \forall j \neq i (w \not\models_{\langle W, D, \phi, \pi, P^* \rangle} b_j))$$

$$\text{or } (w \not\models_{\langle W, D, \phi, \pi, P^* \rangle} a \text{ and } \forall j (w \not\models_{\langle W, D, \phi, \pi, P^* \rangle} b_j))$$

The fourth rule defines the “disjoint” assertion.

$$\begin{aligned} w \models_{\langle W, D, \phi, \pi, P^* \rangle} \text{disjoint}([h_1 : p_1, \dots, h_m : p_m]) \\ \text{if } \exists i (w \models_{\langle W, D, \phi, \pi, P^* \rangle} h_i) \text{ and } \forall j \neq i (w \not\models_{\langle W, D, \phi, \pi, P^* \rangle} h_j) \\ \text{and } \forall i P^*(h_i) = p_i \end{aligned}$$

**Definition A.3** The set of possible worlds given theory  $T$ , denoted  $W_T$  is defined as

$$W_T = \{w \in W : w \models_{\langle W, D, \phi, \pi, P^* \rangle} T'\}$$

where  $T'$  is the set of ground instances of elements of  $T$ .

In order to interpret the probabilities of possible worlds, we create the algebra of subsets of  $W_T$  that can be described by finite formulae of instances of hypotheses.

**Definition A.4**

$$\Omega_{\langle W, D, \phi, \pi, P^* \rangle} = \{\omega \subseteq W_T : \exists \text{formula } f_\omega, \forall w \in \omega w \models_{\langle W, D, \phi, \pi, P^* \rangle} f_\omega\}$$

By formula we mean a finite formula make up of conjunctions and disjunctions of elements of  $H'$ .

The elements of  $\Omega_{\langle W, D, \phi, \pi, P^* \rangle}$  form a sample space [2]. Elements of  $\Omega_{\langle W, D, \phi, \pi, P^* \rangle}$  are closed under finite union, and complementation (given that we can complement any hypothesis by using the disjunct of the remaining hypotheses). Like Bacchus [2], we do not require sigma-additivity of our sample space. Because our language is so weak, we do not need countable unions.

**Definition A.5** An interpretable DNF formula of hypotheses is a formula of the form

$$\bigvee_{i=1}^n \bigwedge_{j=1}^{k_i} h_{ij}$$

where  $h_{ij}$  is a ground instance of a hypothesis such that:

1. for each  $i$ , there is no  $j_1 \neq j_2$  such that  $\{h_{ij_1}, h_{ij_2}\}$  is a subset of an instance of a disjoint declaration in  $T$  (i.e.,  $\{h_{ij_1}, h_{ij_2}\}$  does not form part of an integrity constraint).

2. for each  $i_1 \neq i_2$  there exists  $j_1, j_2$  such that  $\{h_{i_1 j_1}, h_{i_2 j_2}\}$  are in an instance of a disjoint declaration in  $T$  (i.e., the disjuncts are disjoint).
3. For no  $i_1 \neq i_2$  is it the case that  $\{h_{i_1 1}, \dots, h_{i_1 k_{i_1}}\} \subseteq \{h_{i_2 1}, \dots, h_{i_2 k_{i_2}}\}$ .

**Lemma A.6** *Every finite formula make up of conjunctions and disjunctions of elements of  $H'$  is equivalent, given  $T$  (i.e. they describe the same subset of  $W_T$ ) to an interpretable DNF formula.*

**Proof:** To satisfy the first and third conditions we can remove inconsistent conjuncts and any supersets of other formulae. The resulting formula is equivalent to the original.

Suppose  $C_1$  and  $C_2$  are conjuncts that do not satisfy the second condition. Suppose, without loss of generality that  $C_1 = \bigwedge_{j=1}^{k_1} h_{1j}$ . Each  $h_{1j}$  is in a disjoint declaration. Let  $D_j$  be the disjunct of the hypotheses in the instance of a disjoint declaration in which  $h_{1j}$  appears. Each  $D_j$  is true in all elements of  $W_T$ . Thus,  $C_2 \wedge \bigwedge_{j=1}^{k_1} D_j$  is equivalent to  $C_2$  given  $T$ . Distribute  $C_2 \wedge \bigwedge_{j=1}^{k_1} D_j$  into DNF, remove the conjunct that is a superset of  $C_1$ , replace  $C_2$  by this disjunct. We now have removed a violation of the second condition and have an equivalent formula. This can be done for all violations of the second formula.  $\square$

We can define the measure elements of  $\Omega_{\langle W, D, \phi, \pi, P^* \rangle}$  by

**Definition A.7** Suppose  $\omega \in \Omega_{\langle W, D, \phi, \pi, P^* \rangle}$  and  $\omega$  can be described by formula  $f_\omega$ . Suppose  $f_\omega$  is equivalent to interpretable DNF formula  $\bigvee_i \bigwedge_j h_{ij}$ , then define

$$\mu_{\langle W, D, \phi, \pi, P^* \rangle}(\omega) = \sum_i \prod_j P^*(h_{ij})$$

**Lemma A.8**  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}$  is a well-defined probability function, obeying the axioms of probability:

1.  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}(\omega) \geq 0$  for all  $\omega \in \Omega_{\langle W, D, \phi, \pi, P^* \rangle}$
2.  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}(W_T) = 1$ .
3. If  $\omega_1 \cap \omega_2 = \{\}$  then  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}(\omega_1 \cup \omega_2) = \mu_{\langle W, D, \phi, \pi, P^* \rangle}(\omega_1) + \mu_{\langle W, D, \phi, \pi, P^* \rangle}(\omega_2)$ , for all  $\omega_1, \omega_2 \in \Omega_{\langle W, D, \phi, \pi, P^* \rangle}$

**Proof:** To show that it is well defined, we show that any two equivalent formulae give the same value. Given two formulae, consider the set of all instances of hypotheses in the formulae. This subset of  $H'$  is finite. Each hypotheses in this set appears in a ground instance of a disjunct declaration. Let  $D$  be the set of all such disjunct declarations. Each disjunct declaration corresponds to a disjunct of the hypotheses of the disjunct declaration. Each conjunct has elements from a subset  $D$ . We can extend each conjunct to cover  $D$ , by conjoining to each hypothesis the disjunct of each disjunct declaration that is not represented in the conjunction. We then distribute to DNF. This procedure does not change the probability of any conjunct (as the probabilities of the disjunct of each disjunct declaration sum to one, and we can distribute multiplication over addition).

Once this procedure is carried to the two equivalent formulae, they will be syntactically identical (up to commutativity and associativity), as any difference can be extended into a possible world in which they have a different value. This cannot happen as they are equivalent.

Thus they must have the same measure.

It is straightforward to check that the axioms of probability hold for  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}$ .  $\square$

**Definition A.9** If  $T$  is a probabilistic Horn abduction theory, and  $a$  is a formula, the probability of  $a$  given  $T$ , written  $P_T(a)$  is defined as:

$$P_T(a) = \mu_{\langle W, D, \phi, \pi, P^* \rangle}(\{w \in W_T : w \models_{\langle W, D, \phi, \pi, P \rangle} a\})$$

The following lemma shows that  $P_T$  does indeed coincide with  $P^*$ .

**Lemma A.10** If  $h$  is a ground hypothesis defined in  $T$  then  $P_T(h) = P^*(h)$ .

**Proof:**

$$P_T(h) = \mu_{\langle W, D, \phi, \pi, P^* \rangle}(\{w \in W_T : w \models_{\langle W, D, \phi, \pi, P \rangle} h\})$$

$\{w \in W_T : w \models_{\langle W, D, \phi, \pi, P \rangle} h\}$  is described by  $h$ , and so, by definition of  $\mu_{\langle W, D, \phi, \pi, P^* \rangle}$ ,  $P_T(h) = P^*(h)$ .  $\square$

**Theorem A.11** *If  $A$  is a ground atom or conjunction of ground atoms,*

$$P_T(A) = \sum_{e_i \in \text{expl}(A, T)} \prod_{h_{ij} \in e_i} P(h_{ij})$$

*where  $\text{expl}(A, T)$  is the set of minimal explanations of  $A$  from theory  $T$ .*

In other words, the semantics justifies our use of summing the explanations to find the prior probability of a proposition.

**Proof:** For this proof we treat  $A$  as a set of atoms as well as well as the conjunction of these elements. What is meant should be clear from context.

Based on the acyclicity of  $T$ , we define a well founded ordering over sets of occurrences of atoms in  $T'$ . Because  $T'$  is acyclic (assumption 2.8), there is an assignment of natural number to occurrences atoms in  $T'$  such that the elements of the body of a rule are less than the head of the rule. Call this number the *depth* of the rule. Because there are no rules with a hypothesis as head (assumption 2.7), we can consistently assume that all hypotheses have depth zero.

This ordering is based on the lexicographic ordering of pairs  $\langle d, n \rangle$  where  $d$  is the depth of the element of the set with maximal depth, and  $n$  is the number of elements of this depth. Each time through the recursion either  $d$  is reduced or  $d$  is kept the same and  $n$  is reduced. This is well founded as both  $d$  and  $n$  are non-negative integers.

For the base case, where  $d = 0$ ,  $A$  is a conjunction of hypotheses. If  $A$  is inconsistent then there are no explanations, and the models of  $W_T$  can be described by *false*, the empty disjunct, and so  $P_T(A) = 0$ . If  $A$  is consistent, then there is one explanation for  $A$ , namely  $A$  itself, so

$$\sum_{e_i \in \text{expl}(A, T)} \prod_{h_{ij} \in e_i} P(h_{ij}) = \prod_{h \in A} P(A)$$

Because the space defined by  $A$  has  $A$  as its interpretable DNF formula, we have

$$P_T(A) = \prod_{h \in A} P(A).$$

For the inductive case, suppose  $d > 0$  and  $a$  is a proposition in  $A$  with greatest depth. Let  $R = A \setminus \{a\}$ . As  $d > 0$ ,  $a$  is not a hypothesis. Suppose there are  $m$  rules defining  $a$ :

$$\begin{aligned} a &\leftarrow B_1 \\ a &\leftarrow B_2 \\ &\vdots \\ a &\leftarrow B_m \end{aligned}$$

Under the ordering above  $B_i \cup R < d$ , and so we can inductively assume the lemma for  $B_i \cup R$ . Thus

$$P_T(B_i \cup R) = \sum_{e_i \in \text{expl}(B_i \cup R, T)} \prod_{h_{ij} \in e_i} P(h_{ij})$$

Now the explanations of  $a$  are obtained from the explanations of the  $B_i$ ; in particular

$$\text{expl}(\{a\} \cup R, T) = \bigcup_{i=1}^m \text{expl}(B_i \cup R, T)$$

Also, as far as the semantics are concerned  $a \equiv \bigvee_i B_i$ . Thus, by additivity (lemma A.8), and the fact that the  $B_i$  are disjoint

$$\begin{aligned} P_T(\{a\} \cup R) &= \sum_{i=1}^m P_T(B_i \cup R, T) \\ &= \sum_{i=1}^m \sum_{e_i \in \text{expl}(B_i, T)} \prod_{h_{ij} \in e_i} P(h_{ij}) \\ &= \sum_{e_i \in \bigcup_{i=1}^m \text{expl}(B_i \cup R, T)} \prod_{h_{ij} \in e_i} P(h_{ij}) \\ &= \sum_{e_i \in \text{expl}(A, T)} \prod_{h_{ij} \in e_i} P(h_{ij}) \end{aligned}$$

□

## Acknowledgements

Thanks to Alan Mackworth, Andrew Csinger, Michael Horsch, Keiji Kanazawa and George Tsiknis for valuable comments on this paper. Thanks to Mark Wallace for showing me that the language in [43] was too strong. Thanks to Fahiem Bacchus for help with the semantics. This research was supported under NSERC grant OGPOO44121, and under Project B5 of the Institute for Robotics and Intelligent Systems.

## References

- [1] K. R. Apt and M. Bezem. Acyclic programs (extended abstract). In *Logic Programming: Proceedings of the Seventh International Conference*, pages 617–633. MIT Press, 1990.
- [2] F. Bacchus. Lp, a logic for representing and reasoning with statistical knowledge. *Computational Intelligence*, 6(4):209–231, November 1990.
- [3] F. Bacchus. *Representing and Reasoning with Uncertain Knowledge*. MIT Press, Cambridge, Massachusetts, 1990.
- [4] R. Barbuti, P. Mancarella, D. Pedreschi, and F. Turini. A transformational approach to negation in logic programming. *Journal of Logic Programming*, 8:201–228, 1990.
- [5] J. S. Breese. Construction of belief and decision networks. Technical Memorandum 30, Rockwell Palo Alto Lab, Palo Alto, Cal., 1989.
- [6] E. Charniak and R. Goldman. A semantics for probabilistic quantifier-free first-order languages, with particular application to story understanding. In *Proc. 11th International Joint Conf. on Artificial Intelligence*, pages 1074–1079, Detroit, Mich., August 1989.
- [7] E. Charniak and S. E. Shimony. Probabilistic semantics for cost based abduction. In *Proc. 8th National Conference on Artificial Intelligence*, pages 106–111, Boston, July 1990.
- [8] K. L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Databases*, pages 293–322. Plenum Press, New York, 1978.

- [9] L. Console, D. Theseider Dupre, and P. Torasso. On the relationship between abduction and deduction. *Journal of Logic and Computation*, ??(??):??-??, 1991.
- [10] P. T. Cox and T. Pietrzykowski. General diagnosis by abductive inference. Technical Report CS8701, Computer Science, Technical University of Nova Scotia, Halifax, April 1987.
- [11] J. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28(2):127-162, March 1986.
- [12] J. de Kleer, A. K. Mackworth, and R. Reiter. Characterizing diagnoses. In *Proc. 8th National Conference on Artificial Intelligence*, pages 324-330, Boston, July 1990.
- [13] J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97-130, April 1987.
- [14] J. de Kleer and B. C. Williams. Diagnosis with behavioral modes. In *Proc. 11th International Joint Conf. on Artificial Intelligence*, pages 1324-1330, Detroit, August 1989.
- [15] J. Doyle. Methodological simplicity in expert system construction: the case for judgements and reasoned assumptions. *The AI Magazine*, pages 39-43, Summer 1983.
- [16] M. R. Genesereth. The use of design descriptions in automated diagnosis. *Artificial Intelligence*, 24(1-3):411-436, December 1984.
- [17] R. Goebel, K. Furukawa, and D. Poole. Using definite clauses and integrity constraints as the basis for a theory formation approach to diagnostic reasoning. In E. Shapiro, editor, *Proc. Third International Conference on Logic Programming*, pages 211-222, London, July 1986.
- [18] S. Hanks and D. V. McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379-412, 1987.
- [19] J. R. Hobbs, M. E. Stickel, P. Martin, and D. Edwards. Interpretation as abduction. In *Proc. 26th Annual Meeting of the Association for Computational Linguistics*, pages 95-103, Buffalo, June 1988.



- [20] M. Horsch and D. Poole. A dynamic approach to probabilistic inference using Bayesian networks. In *Proc. Sixth Conference on Uncertainty in AI*, pages 155–161, Boston, July 1990.
- [21] H. A. Kautz. The logic of persistence. In *Proc. 5th National Conference on Artificial Intelligence*, pages 401–405, Philadelphia, Penn., August 1986.
- [22] K. Konolige. Closure + minimization implies abduction. technical report ??, SRI International, Menlo Park, CA, 1990.
- [23] R. Kowalski. *Logic for Problem Solving*. Artificial Intelligence Series. North Holland, New York, 1979.
- [24] K. B. Laskey and P. E. Lehner. Assumptions, beliefs and probabilities. *Artificial Intelligence*, 41(1):65–77, 1989.
- [25] R. S. Ledley and L. B. Lusted. Reasoning foundations of medical diagnosis. *Science*, 130(3366):9–21, July 1959.
- [26] F. Lin and Y. Shoham. Argument systems: a uniform basis for non-monotonic reasoning. In *Proc. First International Conf. on Principles of Knowledge Representation and Reasoning*, pages 245–255, Toronto, May 1989.
- [27] J. W. Lloyd. *Foundations of Logic Programming*. Symbolic Computation Series. Springer-Verlag, Berlin, second edition, 1987.
- [28] R. P. Loui. Defeat among arguments: a system for defeasible inference. *Computational Intelligence*, 3(2):100–106, May 1987.
- [29] R. P. Loui. Defeasible decisions: what the proposal is and isn't. In M. Henrion et. al., editor, *Uncertainty in Artificial Intelligence 5*, pages 99–116. North Holland, 1990.
- [30] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In M. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

- [31] E. M. Neufeld and D. Poole. Towards solving the multiple extension problem: combining defaults and probabilities. In *Proc. Third Workshop on Reasoning with Uncertainty*, pages 305–312, Seattle, July 1987.
- [32] J. Pearl. Distributed revision of composite beliefs. *Artificial Intelligence*, 33(2):173–215, October 1987.
- [33] J. Pearl. Embracing causation in default reasoning. *Artificial Intelligence*, 35(2):259–271, 1988.
- [34] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [35] Y. Peng and J. A. Reggia. *Abductive Inference Models for Diagnostic Problem-Solving*. Symbolic Computation – AI Series. Springer-Verlag, New York, 1990.
- [36] J. L. Pollock. Defeasible reasoning. *Cognitive Science*, 11:481–518, 1987.
- [37] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1):27–47, 1987.
- [38] D. Poole. Representing knowledge for logic-based diagnosis. In *International Conference on Fifth Generation Computing Systems*, pages 1282–1290, Tokyo, Japan, November 1988.
- [39] D. Poole. Explanation and prediction: an architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.
- [40] D. Poole. Normality and faults in logic-based diagnosis. In *Proc. 11th International Joint Conf. on Artificial Intelligence*, pages 1304–1310, Detroit, August 1989.
- [41] D. Poole. A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5(5):521–548, December 1990.
- [42] D. Poole. Representing Bayesian networks within probabilistic Horn abduction. In *Proc. Seventh Conf. on Uncertainty in Artificial Intelligence*, pages 271–278, Los Angeles, July 1991.

- [43] D. Poole. Representing diagnostic knowledge for probabilistic Horn abduction. In *Proc. 12th International Joint Conf. on Artificial Intelligence*, pages 1129–1135, Sydney, August 1991.
- [44] D. Poole. Search-based implementations of probabilistic Horn abduction. Technical report, Department of Computer Science, University of British Columbia, Vancouver, B.C., Canada, 1991.
- [45] D. Poole, R. Goebel, and R. Aleliunas. Theorist: A logical reasoning system for defaults and diagnosis. In N. Cercone and G. McCalla, editors, *The Knowledge Frontier: Essays in the Representation of Knowledge*, pages 331–352. Springer-Verlag, New York, NY, 1987.
- [46] D. Poole and G. Provan. What is an optimal diagnosis? In *Proc. Sixth Conference on Uncertainty in AI*, pages 46–53, Boston, July 1990.
- [47] H. E. Pople, Jr. On the mechanization of abductive logic. In *Proc. 3rd International Joint Conf. on Artificial Intelligence*, pages 147–152, Stanford, August 1973.
- [48] G. Provan. An analysis of ATMS-based techniques for computing Dempster-Shafer belief functions. In *Proc. 11th International Joint Conf. on Artificial Intelligence*, pages 1115–1120, Detroit, August 1989.
- [49] H. Reichenbach. *The Direction of Time*. University of California Press, Berkeley and Los Angeles, 1956.
- [50] R. Reiter. Equality and domain closure in first order data bases. *J. ACM*, 27:235–249, 1980.
- [51] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, April 1987.
- [52] R. Reiter and J. de Kleer. Foundations of assumption-based truth maintenance systems: preliminary report. In *Proc. 6th National Conference on Artificial Intelligence*, pages 183–188, Seattle, July 1987.
- [53] R. Reiter and A. K. Mackworth. A logical framework for depiction and image interpretation. *Artificial Intelligence*, 42(2):125–155, 1989.

- [54] R. D. Shachter and D. Heckerman. Thinking backwards for knowledge acquisition. *AI Magazine*, 8:55-62, 1987.
- [55] M. Shanahan. Prediction is deduction, but explanation is abduction. In *Proc. 11th International Joint Conf. on Artificial Intelligence*, pages 1055-1060, Detroit, Mich., August 1989.
- [56] S. E. Shimony and E. Charniak. A new algorithm for finding map assignments to belief networks. In *Proc. Sixth Conf. on Uncertainty in Artificial Intelligence*, pages 98-103, Cambridge, Mass., July 1990.